

5878.P011X

UNITED STATES PATENT APPLICATION
FOR

Tunneling Protocols For Wireless Communications

INVENTOR:

HARRY BIMS

EUGENE LEUNG

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN
12400 WILSHIRE BOULEVARD
SEVENTH FLOOR
LOS ANGELES, CA 90025-1026

(408) 720-8300

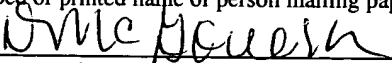
EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number: EV336583726US

Date of Deposit: September 12, 2003

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Assistant Commissioner for Patents, Washington, D. C. 20231

Deborah A. McGovern
(Typed or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)

September 12, 2003
(Date signed)

TUNNELING PROTOCOLS FOR WIRELESS COMMUNICATIONS

RELATED APPLICATIONS

The present patent application is a continuation-in-part (CIP) of, and claims priority to U.S. patent application serial No. 10/044,016, filed on January 11, 2002, which is hereby incorporated reference.

COPYRIGHT NOTICES

[0001] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

FIELD OF THE INVENTION

[0002] The present invention relates to the field of wireless communications; more particularly, the present invention relates to tunneling communication protocols in a wireless communication system.

BACKGROUND OF THE INVENTION

[0003] Figure 1 illustrates an exemplary network environment used today. Referring to Figure 1, a corporate Local Area Network (LAN) backbone 102 interfaces to a number of desktop computers 103_1-103_n and may interface to Internet 101. Corporate LAN backbone 102 may comprise a firewall 102A, corporate server 102B, and a standard Ethernet switch 102C. Ethernet switch 102C includes an interface by which desktops 103_1-103_n are coupled to the corporate LAN backbone 102 for access to corporate sever 102B and to Internet 101 (via firewall 102A).

[0004] More recently, wireless LANs (WLANs) are being installed. Many of the recently implemented WLANs operate according to the protocol set forth in the 802.11 Standard, particularly as more enterprises are adopting the 802.11 Standard. ISO | IEC DIS 8802.11

[0005] Figure 2 illustrates one embodiment of an 802.11 based WLAN (LAN) system. Referring to Figure 2, the Internet or other LAN 201 is coupled to an 802.11 server 203 via firewall (FW) 202. Server 203 communicates with mobile stations in a number of 802.11 cells 206_1-206_n using an access point in each of cells 206_1-206_n , such as access point 204. Server 203 is coupled to access points such as access point 204, via an Ethernet connection. There is one access point for each of the 802.11 cells 206_1-206_n . Mobile stations in each of the 802.11 cells, such as laptops 205_1 and 205_2 in cell 206_1 , communicate wirelessly with the access points via the 802.11 protocol. The communications from mobile stations in the 802.11 cells to the access points are forwarded through to server 203 and potentially to Internet/LAN 201, while communications from Internet/LAN 201 are forwarded through server 203 to the mobile stations via the access points.

[0006] There are a number of problems associated with the current implementations of 802.11 networks. For example, in order to set up an 802.11 network such as shown in Figure 2, a site survey is required in order to determine where each of the access points are to be

placed to ensure that the 802.11 cells provide complete coverage over a particular geographic area. This may be costly. Also, the cost of each of the access points is approximately \$500.00. Generally, such a high cost is a deterrent to having a large number of access points. However, by reducing the number of access points, coverage diminishes and the 802.11 network is less effective. Furthermore, there are a number of mobility problems associated with the current 802.11 network deployments. For example, the 802.11 standard sets forth a number of solutions to handle the issue of mobility of mobile stations between the 802.11 cells. However, these schemes do not work effectively as there is no standard solution in place and users haven't indicated a desire for long-term proprietary solutions.

SUMMARY OF THE INVENTION

[0007] A method and apparatus for tunneling protocols are described. In one embodiment, the method comprises broadcasting a message at a repeater to one or more members in a network including a switch, the broadcasted message indicating that the repeater is entering the network, receiving VLAN (virtual local area network) configuration information from the switch in response to the broadcast message, and communicating with the switch using the VLAN configuration information in subsequent communications.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the invention, which, however, should not be taken to limit the invention to the specific embodiments, but are for explanation and understanding only.

[0009] Figure 1 illustrates an exemplary network environment used today.

[0010] Figure 2 illustrates one embodiment of an 802.11 based wireless LAN-based (LAN) system.

[0011] Figure 3 illustrates one embodiment of a network architecture.

[0012] Figure 4A is a flow diagram of one embodiment of a receiver diversity processing performed by a repeater.

[0013] Figure 4B is a flow diagram of one embodiment of a receiver diversity processing performed by a switch.

[0014] Figure 4C is a process for managing repeaters using a token-based mechanism.

[0015] Figure 4D is one embodiment of a token-based process for handling packets.

[0016] Figure 5A illustrates one technique for location tracking by RSSI.

[0017] Figure 5B is a flow diagram of one embodiment of a process for performing location tracking by a switch.

[0018] Figure 6 illustrates mobility supported by routing.

[0019] Figure 7 illustrates one embodiment of a network system.

[0020] Figure 8 illustrates one embodiment of a protocol architecture.

[0021] Figure 9A illustrates one embodiment of a switch.

[0022] Figure 9B illustrates one embodiment of a repeater.

[0023] Figure 10 illustrates one embodiment of a hardware architecture for a repeater.

- [0024] Figure 11 is a block diagram of one embodiment of the base station processor of a repeater.
- [0025] Figure 12A is a block diagram of one embodiment of a switch.
- [0026] Figure 12B is a flow diagram of one embodiment of a process for reconfiguring the wireless communication system.
- [0027] Figure 13 is one embodiment of a distributed MAC architecture.
- [0028] Figure 14 illustrates one embodiment of the switching plane.
- [0029] Figure 15 illustrates the communication network and exemplary data traffic process.
- [0030] Figure 16 illustrates an exemplary process for transferring data traffic from a mobile station to a desktop.
- [0031] Figure 17 illustrates an exemplary process for transferring data traffic between two mobile stations.
- [0032] Figure 18 illustrates an exemplary process for transferring data traffic from a desktop to a mobile station.
- [0033] Figure 19 is a data flow diagram of one embodiment of an association and token assignment process.
- [0034] Figure 20 is a block diagram of two MAC sublayer instances in a switch.
- [0035] Figure 21 is a data flow diagram of one embodiment of a re-association process.
- [0036] Figure 22 is a flow diagram on one embodiment of a disassociation process.
- [0037] Figure 23A is a block diagram of one embodiment of a communication system.
- [0038] Figure 23B is a block diagram of one embodiment of a communication system having tunneling protocols.
- [0039] Figure 24A is a block diagram of an Ethernet packet.
- [0040] Figure 24B is a block diagram of one embodiment of tunneling protocol packet.

- [0041] Figure 24C is a block diagram of one embodiment of VLAN tag used in a tunneling protocol packet.
- [0042] Figure 24D is a block diagram of one embodiment of tunneling protocol header.
- [0043] Figure 25A is a block diagram of one embodiment of data message Ids used in tunneling protocol packets.
- [0044] Figures 25B and 25C are block diagrams of one embodiment of tunneling protocol packet and its data payload.
- [0045] Figures 26A and 26B are block diagrams of one embodiment of data identifiers used in a payload of a tunneling protocol message.
- [0046] Figures 27A and 27B are block diagrams of embodiments of repeater discovery processes.
- [0047] Figure 28 is a flow diagram of an embodiment of a process for downloading software from a switch to a repeater.
- [0048] Figure 29 is a flow diagram of an embodiment of a process for initially assigning a token to a repeater.
- [0049] Figures 30A and 30B are flow diagrams of embodiments of processes for moving a token from one repeater to another.
- [0050] Figures 31A and 31B are flow diagrams of embodiments of processes for detecting a link down between a switch and a repeater.
- [0051] Figure 32 is a flow diagram of an embodiment of a process for handling a beacon message.
- [0052] Figures 33A and 33B are flow diagrams of embodiments of a processing for data fragmentation.
- [0053] Figures 34A and 34B are block diagrams of exemplary heartbeat messages for a repeater and a switch.
- [0054] Figure 35 is a block diagram of an embodiment of operating states of a repeater.

[0055] Figures 36A and 36B are block diagrams of an embodiment of a set data message and a corresponding response message.

[0056] Figure 37 is a block diagram of one embodiment of VLAN configuration information.

[0057] Figure 38 is a block diagram of one embodiment of an acknowledgment frame for a repeater.

[0058] Figures 39A-39C are block diagrams of one embodiment of an inbound 802.11 management, control, and data frames.

[0059] Figures 39D-39F are block diagrams of one embodiment of an outbound 802.11 management, control, and data frames.

[0060] Figures 40A-40C are block diagram of one embodiment of a token assignment, token deletion, and token query frames.

[0061] Figure 41 is a block diagram of one embodiment of a beacon message.

[0062] Figure 42 is a block diagram of one embodiment of an RSSI information message.

DETAILED DESCRIPTION

[0063] A communication system is described. In one embodiment, the communication system comprises a mobile station having a transmitter to transmit packets wirelessly according to a protocol and multiple repeaters communicably coupled with the mobile station. Each of the repeaters receives one or more packets of the wirelessly transmitted packets from the mobile station. In one embodiment, the repeater assigned to the mobile station (e.g., the repeater with the token of the mobile station) sends the body of each of the one or more packets along with a received signal strength for each of the one or more packets to a switch communicatively coupled to the repeaters. Each of the other repeaters determines which of the wirelessly transmitted packets they received without errors and a received signal strength for those packets. Each of these other repeaters forwards to the switch the received signal strength indication. Based on the received signal strength indication received from each of the repeaters, the switch determines whether to reassign the mobile station to another repeater, such as, for example, a repeater that receives the mobile station packets at a higher received signal strength than any other repeater.

[0064] In one embodiment, the repeaters are grouped and the switch handles each group of repeaters separately. Even so, if a mobile station moves to a location in which a different repeater in a different group is associated with the mobile station, any data buffered by the switch may be forwarded to the mobile device through the new repeater using a single data transfer within the switch.

[0065] In the following description, numerous details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the present invention.

[0066] Some portions of the detailed descriptions which follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0067] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0068] The present invention also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but is not limited to, any type of disk including floppy

disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

[0069] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

[0070] A machine-readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer). For example, a machine-readable medium includes read only memory ("ROM"); random access memory ("RAM"); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.); etc.

Exemplary Network Architecture

[0071] Figure 3 illustrates one embodiment of a network architecture. Referring to Figure 3, a LAN backbone 102 interfaces a number of desktops 103₁-103_n to Internet 101. Note that the present invention does not require that a LAN backbone be included. All that is necessary is that there be a communication mechanism that is capable of receiving packets from other devices and/or sending packets to other devices.

[0072] Similar to Figure 1, LAN backbone 102 includes firewall 102A, corporate server 102B and Ethernet switch 102C. However, in contrast to Figure 1, LAN backbone 102 also includes switch 301 which interfaces to repeaters 302₁-302₃. Although only three repeaters are shown, alternative embodiments may utilize any number of repeaters with a minimum of one. In one embodiment, switch 301 is coupled to repeaters 302₁-302₃ via a wired connection, such as cabling. In one embodiment, the wired connection may comprise CAT5 cabling.

[0073] Each of repeaters 302₁-302₃ receives wireless communications from devices (e.g., mobile stations such as, for example, a mobile phone, a cellular phone, a cordless phone, a headset, a voice-enabled mobile station, a laptop computer system, a personal digital assistant, a computer-data-enabled mobile station, a speakerphone, video game controller, a DVD controller, a stereo controller, a TV controller, etc.) in the coverage areas of the repeaters. In one embodiment, these wireless communications are performed according to the 802.11 protocol. That is, each of the mobile stations in each of cells 310₁-310_n exchanges packets with repeaters 302₁-302₃ using the 802.11 protocol.

[0074] In one embodiment, switch 301 includes 802.11 MAC (medium access control) protocol software and/or hardware that allows switch 301 to communicate with repeaters 302₁-302₃. Different from the prior art, some of the 802.11 MAC functionality typically associated with the access points, as described above in the Background section, are not in repeaters 302₁-302₃ and are, instead, centralized in switch 301. More specifically, the MAC layer is split between repeater 302₁-302₃ and switch 301 to enable transfer of messages over wiring (e.g., CAT5 cabling). As such, repeaters 302₁-302₃ and switch 301 coordinate to perform functionality of the 802.11 MAC layer as described below.

[0075] In one embodiment, switch 301 includes one or more Ethernet connectors (e.g., external Ethernet connector) to enable a computer system, such as desktop computer system 303, or other device, to have an Ethernet connection to LAN backbone 102 via switch 301.

Similarly, in one embodiment, one or more of repeaters 302₁-302₃ includes an Ethernet connector to enable a device (e.g., computer system, such as desktop computer system 304) to gain access, via a repeater, such as repeater 302₃, to switch 301 and the rest of the communication system. In such a case, the wiring coupling switch 301 to repeaters 302₁-302₃ may combine 802.11 information, including management and control (as opposed to solely data) information, with traditional Ethernet packets on the same wiring (e.g., CAT5).

[0076] In one embodiment, switch 301 may be implemented as a server that may be located within corporate LAN backbone 102 or other networks. The server may communicate with repeaters 302₁-302₃ over the network (e.g., a layer 2 network). That is, the server includes, but not limited to, substantially all the functionality of switch 301 to handle communications between the server and repeaters 302₁-302₃ over a layer 2 network. In this embodiment, repeaters 302₁-302₃ are communicatively coupled, wired or wirelessly, to the network where the server is located (e.g., corporate LAN backbone 102), instead of a port of switch 301. The server communicates with the repeaters based on layer 2 information (e.g., MAC layer) of repeaters 302₁-302₃. It will be appreciated that one or more networks (e.g., another Intranet or WAN) may exist between the network of the server (e.g., corporate LAN backbone 102) and repeaters 302₁-302₃.

Distributed Receiver Diversity Approach

[0077] The network architecture described above allows for overlapping coverage between cells supported by the repeaters. This overlapping coverage allows for receiver diversity.

[0078] The packets from the mobile stations in each of the cells are broadcast and may be received by multiple repeaters. By allowing multiple repeaters to receive packets from one of the mobile stations, collisions and dropped packets may be reduced or avoided. For example, if a collision occurs or if a packet is dropped by one of the repeaters, then a particular packet

can still be received by other repeaters. In this manner, the use of repeaters described herein provides for higher reliability.

[0079] In an embodiment in which mobile stations exchange packets with repeaters using the 802.11 protocol, each packet from a mobile station includes an Ethernet MAC address, which is embedded in the packet. Each packet may be received by one or more repeaters. Each repeater that receives a packet from a mobile station without errors (i.e., cleanly) determines the received signal strength of the packet in a manner well-known in the art. The received signal strength is converted into an indication, such as a received signal strength indicator (RSSI). In one embodiment, the RSSI is specified in a value from 1 to 127. These 128 discrete values can be mapped to dB signal strength values based on the particular implementation being used. In one embodiment, all repeaters send their RSSI for the packet to switch 301. The repeater that is assigned to the mobile station (e.g., has the token for the mobile station) sends the packet along with the RSSI. In one embodiment, the repeater encapsulates the packet into an Ethernet packet with the RSSI in a header and forwards the Ethernet packet to switch 301. In an alternative embodiment, each repeater receiving the packet without error forwards the packet, along with the RSSI to the switch. Thus, all packets received from mobile stations by a repeater without errors are forwarded to switch 301. Switch 301 knows which repeater sent the packet(s) because it is received on its preassigned port.

[0080] In one embodiment, the fact that a particular repeater received a packet without errors is communicated to all other repeaters. In one embodiment, this is accomplished by having the repeater send each encapsulated packet and its RSSI as a broadcast packet to switch 301. This broadcast packet is similar to those broadcast packets used in Ethernet and includes a special broadcast address, which is recognized by switch 301. In another embodiment, only the header of the packet, which includes the RSSI and uniquely identifies

the packet, is encapsulated and sent as a broadcast packet to the other repeaters. In this case, the data portion of the packet is not forwarded.

[0081] In response to receiving the broadcast packet with the specific broadcast address, switch 301 broadcasts the packet on all of the other ports used for communication between switch 301 and the other repeaters.

[0082] In one embodiment, upon receiving a packet without error from a particular mobile station, the repeater sets a timer within which it is to receive packets received by other repeaters that are duplicates to the packet it has already received. When the timer expires, the repeater examines the RSSI of the packet it received (without error) with the RSSI values of duplicate packets received by other repeaters. Based on that information, the repeater determines if it is to send an acknowledgement packet for the subsequent transmission from the mobile station. Thus, if the time expires without receiving a duplicate packet, the repeater sends the acknowledgement. If the timer expires and the repeater receives a duplicate packet, thereafter, it is treated as a new packet. To avoid this, the timer time out value is set to handle the worst case time delay that a repeater may face in receiving duplicate packets.

[0083] In one embodiment the communication protocol used between the mobile stations and the repeater is a modified version of the 802.11 protocol in which acknowledgement packets are disabled and thus, not sent. By doing so, a repeater that is assigned to a particular mobile station in response to receiving a packet at a larger received signal strength than any other repeater may not be delayed in handling the packet, and thereby avoids missing a portion of a packet from the mobile station.

[0084] Note that switch 301 forwards each packet received from repeaters (note duplicates) to the remainder of the communication system (e.g., LAN backbone, other mobile stations, the Internet, etc.). In one embodiment, this occurs after de-duplication of packets so that only one copy of each packet is forwarded.

[0085] Once the broadcast packets have been received, in one embodiment, all the repeaters know what packets were received cleanly by the others and at what RSSI the packets were received by the other repeaters. Thereafter, each repeater selects the packet with the highest RSSI and determines the repeater that received it. In other words, each repeater performs a comparison on the received signal strength of the packets it received that were also received by one or more other repeaters. For each of the packets that a repeater receives at a power level higher than any of the other repeaters that received that packet, that repeater sends an acknowledgement back to the mobile station acknowledging that the packet was received without errors. This prevents all the repeaters that receive the packet cleanly from sending multiple acknowledgements to the mobile station.

[0086] In one embodiment, if two repeaters have the same receive signal strength for a packet, the repeater with the lower port number (the port number by which switch 301 is coupled to the repeater) is the repeater that is selected to send the acknowledgement to the mobile station. In this manner, only one repeater is selected to send the acknowledgement to the mobile station and, thus, the receiver diversity is handled in the network architecture in a distributed fashion. In one embodiment, to enable the repeaters to determine which is to send the acknowledgement in the case of a packet received with the same received signal strength by multiple repeaters, each packet includes identification information, such as its switch port number, to enable the determination of which has the lowest port number. Note, in an alternative embodiment, the repeater with the highest port number may be the one to send the acknowledgement or other pre-assigned priority information may be used by the repeaters in such situations. In an alternative embodiment, instead of using the port number to determine which repeater is to send the acknowledgement when two repeaters have the same received signal strength, the repeater MAC address may be used. For example, the repeater with the lowest MAC address may send the acknowledgement packet, or alternatively, the repeater with the highest MAC address may send the acknowledgement packet. Using the repeater

MAC address for the determination is preferred in cases where a layer 2 network is communicatively coupled between the repeaters and the switch such that a single port is used by multiple repeaters.

[0087] Figure 4A is a flow diagram of one embodiment of a receiver diversity process performed by a repeater. The process is performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both.

[0088] Referring to Figure 4A, processing logic initially receives an 802.11 packet (processing block 401). In response to the 802.11 packet, processing logic determines the received signal strength (e.g., RSSI) (processing block 402). In one embodiment, this processing logic comprises a hardware mechanism, such as a radio frequency (RF) device (e.g., integrated circuit (e.g., RF IC 1002 in Figure 10)) in the repeater. In such a case, the RF device sends the RSSI to a baseband processor in the repeater.

[0089] Thereafter, processing logic encapsulates 802.11 packet and RSSI in an Ethernet packet (processing block 403) and sends the Ethernet packet to the switch (processing block 404). In one embodiment, a baseband processor (e.g., baseband processor 1001 in Figure 10) performs the encapsulation and sends the Ethernet packet to the switch.

[0090] Later in time, processing logic receives one or more packets from the switch that are duplicates of the 802.11 packet. These duplicate packets are transmitted by other repeaters and encapsulated by those repeaters, along with their RSSIs (processing block 405).

Processing logic in the repeater compares RSSIs for the duplicate packets (processing block 406). In one embodiment, a baseband processor (e.g., baseband processor 1001 in Figure 10) performs the comparison. If the repeater determines it received the 802.11 packet with the highest RSSI, then processing logic sends the acknowledgment packet to the mobile station (processing block 407).

[0091] Figure 4B is a flow diagram of one embodiment of a receiver diversity processing performed by a switch. The process is performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both.

[0092] Referring to Figure 4B, processing logic initially receives a packet from a repeater (processing block 411). In response to the packet, processing logic determines that the packet is to be sent to the other repeaters and re-broadcasts the received packet to other repeaters (processing block 412). Then processing logic sends only one copy of the packet to the rest of the network (processing block 413).

Token-based Receiver Diversity Approach

[0093] Note that the above receiver diversity procedure is particularly useful when gigabit or faster Ethernet communication exists between switch 301 and repeaters 302₁-302₃. However, if such is not the case, another technique for receiver diversity may be utilized. For example, a token-based receiver diversity procedure may be used. In this case, switch 301 has a token for every mobile station on the 802.11 network and it gives the token to one of the repeaters. In other words, switch 301 pre-assigns the token before a packet is even transmitted by a mobile station. The repeater stores the token in a table that lists all mobile stations for which it has a token. The repeater with the token sends the acknowledgement packet to the mobile stations listed in the table when those mobile stations send packets that are received by the repeater. Therefore, a comparison of received signal strengths for duplicate packets is not necessary. Note that in one embodiment with this token based mechanism, if the repeater with the token does not receive a packet cleanly, but another repeater does, that packet will be forwarded to the switch and not acknowledged to the mobile client. However, switch 301 moves the token before a subsequent packet is sent by the mobile station. Therefore, this will only occur for one packet.

[0094] In one embodiment, switch 301 includes a database with a listing of mobile stations and repeater numbers corresponding to the repeater that has been designated to acknowledge packets received from the mobile station and, thus, has the token. The table may also include additional information describing the repeater itself.

[0095] Since switch 301 receives the received signal strength from each repeater that received a packet without error, switch 301 can determine the closest repeater to a particular mobile station. If the repeater determined to be closest to the particular mobile station is different than the one previously identified as closest (e.g., based on RSSI of packets received from the mobile station), then switch 301 moves the token to a new repeater, i.e. the one that is closer to the mobile station. The token may be moved on a packet-by-packet basis or every predetermined number of the packets (e.g., 10 packets, 100 packets, etc.).

[0096] Switch 301 may employ a timer to indicate the time during which duplicate RSSI values for the same packet may be received in much the same manner the timer is used by the repeaters in the distributed approach described above.

[0097] Figure 4C is a process for managing repeaters using a token-based mechanism. The process is performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both.

[0098] Referring to Figure 4C, processing logic first determines the location of mobile stations with respect to repeaters (processing block 451). Processing logic then assigns a token for each of the mobile stations to one of the repeaters (processing block 452) and stores an indication of the repeater assigned to each mobile station (processing block 453). This information is stored in a table in memory. This table is referred to herein as an active station list. In one embodiment, this table includes a listing of mobile stations and an indication of which repeater and/or switch port number is assigned to the mobile station. The table may be the same data structure used for location tracking described below.

[0099] In one embodiment, the switch assigns a token by sending an Add Token command to the repeater, which causes the repeater to add a new mobile station to its table of mobile stations that the repeater supports. This command includes the MAC address of the mobile station.

[00100] Subsequently, processing logic periodically tests whether the repeater assigned the token for a particular mobile station is still the closest repeater to that mobile station (processing block 454). If so, then the processing is complete. If not, then processing logic moves the token to the closest repeater (processing block 455) and updates the table (e.g., the active station list) to reflect the new repeater that is closest to the mobile station (processing block 456). Processing logic also updates the switch port to reflect the new repeater for use when sending packets to the mobile station from the switch.

[00101] In one embodiment, the switch moves the token by sending a Delete Token command to the repeater that currently has it, causing the repeater to delete the token (and assorted MAC Address) from its list of supported mobile stations, and by sending an Add Token command to the repeater that is currently closest to the mobile station.

[00102] Figure 4D is one embodiment of a token-based process for handling packets. The process is performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both.

[00103] Referring to Figure 4D, processing logic receives a token from the switch (processing block 470) and stores the token in a table stored in a repeater memory that indicates all the mobile stations for which the repeater has a token (processing block 471).

[00104] Subsequently, when processing logic receives a packet from mobile station (processing block 472), processing logic compares the MAC address of the 802.11 packet from the mobile station with the address in the table (processing block 473). Then, processing logic tests whether the MAC address of a packet equals an address in the table (processing

block 474). If so, processing logic provides an acknowledgment (ACK) packet to the mobile station (processing block 475). If not, processing logic ignores the packet.

[00105] Note that since all repeaters communicate the fact that they received a packet from a mobile station along with the received signal strength to switch 301, switch 301 is able to determine the coverage area of the transmission of the mobile station. In one embodiment, each packet received by switch 301 from the repeaters terminates in a network processor in switch 301 (e.g., network processor 1206 of Figure 12), which determines the coverage area because it has access to the RSSI values. By determining the coverage area of the transmission, switch 301 is able to track the location of a particular device.

Downstream Communication Scheduling

[00106] For communications in the reverse direction (e.g., in the downstream direction), in one embodiment, the repeater transmissions are scheduled to reduce collisions. This scheduling is useful because repeaters can be close enough to interfere with one another. Because of this, switch 301 schedules the transmissions to prevent the collisions when the repeaters are actually transmitting.

[00107] For example, if a packet is destined for a particular IP address, then switch 301 performs an address translation to translate, for example, the IP address into an Ethernet MAC address. Switch 301 uses the Ethernet MAC address to search in a location tracking database to determine which repeater is closest to the mobile station having the Ethernet MAC address. Once the repeater is identified by switch 301, then switch 301 knows the switch port on which the packet should be sent so that it is sent to the repeater listed in the location tracking database (for forwarding by the repeater to the mobile station).

[00108] Once the repeater (and the port number) has been identified, switch 301 checks whether an interference problem would be created if the packet is sent by switch 301 to the mobile station at that time. An interference problem would be created if there are other

transmissions that would be occurring when the packet is forwarded onto its destination mobile station. If no interference problem would exist, switch 301 sends the packet through the identified port to the repeater most recently determined to be closest to the mobile station. However, if an interference problem would be created by sending the packet immediately, then switch 301 delays sending the packet through the identified port to the repeater most recently determined to be closest to the mobile station.

[00109] In one embodiment, to determine if an interference problem would exist if a packet is sent immediately upon determining the switch port number on which the packet is to be sent, switch 301 maintains and uses two databases. One of the databases indicates which of the repeaters interfere with each other during their transmissions. This database is examined for every downstream packet that is to be sent and switch 301 schedules the transmission of downstream packets so that the repeaters that interfere with each other when they transmit at the same time do not transmit at the same time. The other database is a listing of mobile stations and the corresponding set of repeaters that last received the transmissions. If two mobile stations have overlapping sets, then it is possible for their acknowledgement packets to interfere when they simultaneously receive non-interfering data packets from different repeaters. Because the mobile stations send acknowledge packets upon receiving downstream packets, there is a possibility that the mobile stations will interfere with each other when sending their acknowledgement packets. Switch 301 takes this information into account during scheduling and schedules downstream packets to the mobile stations to reduce the occurrence of mobile stations interfering with other when sending acknowledgment packets. The information in these two databases may be collected by sending out test packets to the WLAN to determine which repeaters and mobile devices cause the interference described above.

[00110] Alternatively, in one embodiment, all repeaters communicatively coupled to the switch perform the scheduling instead of the switch. This type of scheduling is also referred

to as a distributed form of scheduling, which is in addition to the centralized scheduling typically performed by a switch. In one embodiment, there is no need to have databases to keep track of where the interference may occur (e.g., between the communications of repeaters and mobile stations). Instead, each repeater acts on its own to transmit packets when the repeater determines that the communication channel is clear. For example, if a repeater detects that there may be one or more transmissions performed by one or more other repeaters that would cause interference that prevents the packets from being received by the intended mobile station, the repeater may wait for a period of time to allow the respective communication channel cleared, using some techniques similar to a CSMA/CD (carrier sense multiple access/collision detection) algorithm. Alternatively, the repeaters may communicate with each other, wired or wirelessly, with respect to the scheduling using, for example, tunneling protocols within the Ethernet protocol. However, in one embodiment, in such a case, the switch still handles routing the token to the correct repeater for mobility support.

Upstream Communication Scheduling

[00111] The same databases used for downstream traffic scheduling may be used for upstream traffic scheduling to enable the switch to schedule upstream communications. As with downstream traffic, by using the two databases, the switch is able to determine when parallel communication may take place based on the overlap described above. This scheduling may be used to implement quality of service (QoS) where a period of time when traffic is scheduled is used (bypassing the CSMA algorithm). Similarly, as described above, the scheduling may be performed by the repeaters to coordinate the upstream transmissions to the switch without involving the switch.

Location – Tracking by Received Signal Strength (RSSI)

[00112] Figure 5A illustrates one technique for location tracking by RSSI. Referring to Figure 5A, switch 301 obtains the RSSI for each packet received by the repeaters and may have multiple RSSI values for a packet when that packet is received by two or more different repeaters. More specifically, a mobile station communicates with two (or more) repeaters and one repeater is going to have a stronger received signal strength than the other for the same packet. Based on this information, switch 301 is able to determine that a mobile station is closer to one repeater than the other. By continually monitoring the received signal strength, switch 301 can track the movement of a mobile station with respect to the repeaters.

[00113] Figure 5B is a flow diagram of one embodiment of a process for performing location tracking by a switch. The process is performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. In one embodiment, the processing logic comprises a network processor in the switch (e.g., network processor 1206 of Figure 12).

[00114] Referring to Figure 5B, processing logic compares the RSSI for the duplicate packets received by different repeaters from a mobile station (processing block 550) and tests whether the repeater with the highest RSSI for the packet is the repeater listed as closest to the mobile station in a location tracking table (e.g., database) (processing block 551). If not, processing logic updates the table to indicate that the repeater that received the packet with the highest RSSI is the closest repeater (processing block 552). Processing logic also switches port assignment for the mobile station to the port of new repeater (if the port is different than the port of the previously assigned repeater).

[00115] In one embodiment, the location tracking table may include a listing of mobile stations and their individually assigned repeaters. The location tracking table may also be referred to herein as the active station list. This table may also include, or include instead of

the assigned repeater, an indication of the switch port by which the switch is to communicate with the repeater assigned to each mobile station.

Mobility Supported By Routing

[00116] Figure 6 illustrates mobility supported by routing. Referring to Figure 6, the dotted arrow path for communication from switch 301 to mobile station 601 through repeater 302₂ is the original communication path with the network. As the mobile station 601 moves, a routing handoff occurs so that communication occurs over the solid arrowed path. In order to accomplish this handoff, switch 301 reroutes the packet to a different port. For example, if the first communication path illustrated as the dotted line arrow was on port 1, switch 301 may switch the packet to port 5, the port that associated with the communication path through repeater 302₁. Thus, mobility is supported by simply moving a packet to a different port of switch 301 that is assigned to a different repeater. In such a situation, the mobility provisions of the 802.11 protocol may be ignored. Note that for embodiments where multiple repeaters are on the same port, the repeater MAC address may be changed instead of changing the port number.

[00117] In one embodiment, switch 301 determines that a particular mobile station is closer to a different repeater (by monitoring the received signal strength of duplicate packets). As described above, switch 301 maintains a table (e.g., database, active station list, etc.) of all mobile stations in the 802.11 network and includes an indication of the repeater closest to each mobile station. Switch 301 performs port-based routing and may use the table in the same manner an IP routing table is used. Switch 301 has an Ethernet port for each repeater. When switch 301 determines that a mobile station is closer to a repeater that is different than the one listed in the database (based on the received signal strength of duplicate packets among multiple repeaters), then switch 301 updates the database. Thereafter, if a packet is received by switch 301 for that mobile station, switch 301 merely sends it out on the Ethernet

port assigned to the repeater that was most recently determined to be the closest to that mobile station.

Multi-Switch System

[00118] Figure 7 illustrates one embodiment of a multi-switch system. Referring to Figure 7, the network architecture includes switches 701 and 702 are communicably coupled to server 712. In one embodiment, server 712 is part of a LAN backbone through which access to the Internet and other resources is made. Alternatively, server 712 may act as an interface to another portion of the communication system. Each of switches 701 and 702 is coupled to one or more repeaters in the same manner as described above with respect to Figure 3. In still another embodiment, server 712 may exist within one of, or both, switches 701 and 702.

Protocol Architecture

[00119] Figure 8 illustrates one embodiment of a protocol architecture. Referring to Figure 8, switch 801 is shown having a network layer 801A and a MAC layer 801B. In one embodiment, the network layer 801A comprises a TCP/IP network layer. MAC sublayer 801B communicates with a MAC sublayer of each of repeaters 802₁-802_N. Thus, in contrast to the prior art in which the 802.11 MAC layer is completely within the access point, the 802.11 MAC layer is split between switch 301 and repeaters 802₁-802_N, and the MAC sublayer of the repeaters performs much less functionality than the MAC sublayer of the access points described above.

[00120] In one embodiment, the repeater MAC sublayer is responsible for performing portions of the 802.11 protocol including handling CSMA/CA, DIFS/EIFS interframe spacing (IFS) timing, SIFS timing and control, generating acknowledgement (of ACK) frames (during transmit only) on data packets received, such as 802.11 data frames and generating CTS (clear-to-send) frames in response to RTS (request-to-send) frames. The repeater MAC

sublayer may also respond to the resetting of internal network allocation vectors (NAVs) which are embedded into certain frames (e.g., RTS and CTS frames). Each of the above repeater MAC functions may be implemented in a manner that is well-known is the art.

[00121] In addition to the MAC sublayer, each of repeaters 802₁-802_N includes an 802.11 physical layer or other wireless physical layer.

[00122] The switch MAC sublayer is responsible for handling multiple frame types during reception from the repeaters. In one embodiment, the MAC frame types the switch is capable of handling include an association request, reassociation request, probe request, ATIM, disassociation, authentication, deauthentication, PS-Poll, CTS (updates NAV in repeaters), ACK (in response to data frames), data and Null frames.

[00123] The switch MAC frame types that are accommodated during transmission include an association response, a reassociation response, probe response, ATIM (announcement traffic indication message), disassociation, deauthentication, PS-Poll, data frames, Null frames, RTS (updates NAV in repeater), and beacon frames. It should be noted that the MAC frame types that the switch accommodates during receive and transmit are well known in the arts and part of the 802.11 standard. Each of the above switch MAC functions may be implemented in a manner that is well-known is the art.

[00124] Figure 9A illustrates an embodiment of a switch. In one embodiment, exemplary switch 900 includes session management unit 901, protocol unit 902, location tracking unit 903, fragmentation unit 904, DCF (distributed coordination function) unit 905, packet de-duplication unit 906 and SNMP (simple network management protocol) unit 907. Some of these units may be implemented within the corresponding MAC layer of the switch. In one embodiment, session management unit 901, which may be a part of the switch management entity (SwME), is responsible for handling initial handshakes with one or more repeaters and the associated mobile devices, including, but not limited to authentication, association, and disassociation, etc. In one embodiment, protocol unit 902 handles a variety of network

protocols including 802.11 wireless protocol, RADIUS, VPN (virtual private network) protocol, as well as other wireless protocols, such as, for example, 802.15 (wireless personal area network or WPAN, also referred to as Bluetooth) protocol or 802.16 (broadband wireless metropolitan area network) protocol.

[00125] Location tracking unit 903 is responsible for tracking one or more mobile stations communicating with one or more repeaters using one of the aforementioned mechanisms shown in Figures 5A and 5B. In one embodiment, location tracking unit 903 obtains the RSSI for each packet received by the repeaters and may have multiple RSSI values for a packet when that packet is received by two or more different repeaters. More specifically, a mobile station communicates with two (or more) repeaters and one repeater typically has a stronger received signal strength than the other for the same packet. Based on this information, location tracking unit 903 is able to determine that a mobile station is closer to one repeater than the other. By continually monitoring the received signal strength, location tracking unit 903 tracks the movement of a mobile station with respect to the repeaters. Based on this information, location tracking unit 903 may automatically perform an operation similar to a site survey and may reconfigure the communication network, which will be described in details further below. In one embodiment, location tracking unit 903 may notify session management unit 901 to perform such operations. In a further embodiment, location tracking unit 903 may further detect whether an interference between multiple mobile stations exists and may signal session management unit 901 to perform further action, such as, for example, queuing and rescheduling of requests, such that multiple mobile stations may be able to share a communication channel (e.g., single frequency band), which will be described further below.

[00126] Fragmentation unit 904 is responsible for fragmenting packets to improve performance in the presence of RF interference detected using one of the aforementioned processes. The use of fragmentation can increase the reliability of frame transmissions. By

sending smaller frames, collisions are much less likely to occur. The fragment size value may be typically set between 256 and 2,048 bytes. However, this value may be user controllable via, for example, a user interface of session management unit 901.

[00127] DCF unit 905 is responsible for handling DCF functionality according to 802.11 specifications. DCF unit 905 typically operates based on the CSMA/CA (carrier sense multiple access with collision avoidance) protocol. In a conventional approach, typical 802.11 stations contend for access and attempt to send frames when there is no other station transmitting. If another station is sending a frame, the stations wait until the channel is free. According to one embodiment, when DCF unit 905 detects if interference exists between multiple mobile stations, DCF unit 905 may signal session management unit 901 to queue up the frames and reschedule to processing of these frames to avoid interference. As a result, mobile stations do not need to worry about interference and wait for a free channel, which leads to a wider bandwidth of the network.

[00128] As described above, switch 900 may receive multiple identical packets from multiple repeaters because the corresponding mobile station may send the same packet to multiple repeaters within a coverage area. The switch may broadcast the packet to the rest of the repeaters. In order to avoid broadcasting the same packet multiple times, switch 900 may invoke packet de-duplication unit 906 to de-duplicate the duplicated packets and only one of multiple instances of the same packet gets broadcasted.

[00129] SNMP unit 907 performs typical network management operations well known in the art. SNMP is a protocol governing network management and the monitoring of network devices and their functions. It is not necessarily limited to TCP/IP networks.

[00130] Figure 9B illustrates an embodiment of a MAC sublayer of a repeater. In one embodiment, the repeater MAC sublayer 908 is responsible for performing portions of the 802.11 protocol including handling CSMA/CA, DIFS/EIFS interframe spacing (IFS) timing, SIFS timing and control (block 912), generating acknowledgement (of ACK) frames (during

transmit only) on data packets received, such as 802.11 data frames (block 911) and generating CTS (clear-to-send) frames in response to RTS (request-to-send) frames. The repeater MAC sublayer may also respond to the resetting of internal network allocation vectors (NAVs) which are embedded into (e.g., RTS and CTS frames) (block 910). Each of the above repeater MAC functions may be implemented in a manner that is well-known in the art.

[00131] Figure 10 illustrates one embodiment of a hardware architecture for a repeater. Referring to Figure 10, an RF chip 1002 receives and transmits RF transmissions using antenna 1003. In one embodiment, RF chip 1002 comprises a standard 802.11 RF chip. In one embodiment, antenna 1003 comprises a dual-diversity antenna. Communications received by RF chip 1002 are forwarded on to baseband processor 1001, which is a digital chip that is described in further detail below. Similarly, transmissions to be sent are received by RF chip 1002 from baseband processor 1001.

[00132] Baseband processor 1001 is a digital chip that performs the reduced MAC functions as described above. The repeater also includes a port 1007 for coupling to a switch, such as switch 301. Baseband processor 1001 handles communication with switch 301 using port 1007. In one embodiment, port 1007 also transfers information (through the port) at 100Mb/s bits per second. Port 1007 may also provide power to baseband processor 1001.

[00133] A desktop port 1006 may be included to allow desktop or other systems to plug into the repeater. Also, in one embodiment, LEDs 1005, such as an activity LED, power LED, and/or link LED, may be included in the repeater as well.

[00134] Figure 11 is a block diagram of one embodiment of the baseband processor of a repeater. Baseband processor 1001 includes a repeater MAC and control unit 1105 that interfaces with RF chip 1002 using a protocol. In one embodiment, the interface comprises a TCP/IP layer and an 802.11 MAC sublayer. The repeater MAC/control unit 1105 is coupled to switch 1103. In one embodiment, MAC/control unit 1105 communicates with switch 1103

using a TCP/IP layer and an 802.11 MAC sublayer tunneled inside Ethernet packets. Switch 1103 is also coupled to MAC/PHY layer unit 1104 which interfaces the baseband processor to desktop port 1006. Switch 1103 is also coupled to the activity/power/link LEDs 1005.

Similarly, switch 1103 is coupled to the MAC/physical layer unit 1001 that interfaces the rest of the components on baseband processor 1001 to switch port 1007 via switch 1103. Also coupled to switch port 1007 is a power distribution unit 1102. In one embodiment, power distribution unit 1102 obtains power from the CAT5 wiring and provides it to the rest of baseband processor 1001.

[00135] Figure 12A is a block diagram of one embodiment of a switch. Referring to Figure 12, the switch includes one or more ports 1201 to repeaters 1201. Although 12 are shown, any number may be included. Ports 1201 are coupled to a switching processor 1202. In one embodiment, switching processor 1202 switches 13 ports of gigabit Ethernet and allows broadcast packets to be received on one port and broadcast on the others without involving the rest of the switch. In one embodiment, switching processor 1202 comprises a BCM 5633 gigabit switching processor from Broadcom Corporation of Irvine, California.

[00136] HyperTransport controller 1203 is coupled to switching processor 1202 and provides a gigabit Ethernet interface to the rest of the switch architecture. In one embodiment, the HyperTransport controller 1203 includes a diagnostic port 1204 and another Ethernet port 1205 for use, for example, in coupling to a corporate LAN.

[00137] In one embodiment, HyperTransport controller 1203 comprises a Gallileo HyperTransport controller from Marvell of Sunnyvale, California.

[00138] A network processor 1206 is coupled to HyperTransport controller 1203 and performs the majority of the functions of the switch, including the receiver diversity functions and location-tracking functions described herein, with the exception of the rebroadcast of the broadcast packets received by the switch, which is handled by switching processor 1202. In one embodiment, network processor 1206 is coupled to a boot memory 1209, a DRAM 1207

and one or more LEDs 1208. In one embodiment, network processor 1206 comprises a PMC-Sierra RM9000X2 sold by PMC-Sierra of Santa Clara, California, boot memory 1209 comprises an MB boot flash AMD AM29LV640D boot flash memory and DRAM 1207 comprises 64MB synchronous DRAM (SDRAM) from Advanced Micro Devices, Inc. of Sunnyvale, California.

[00139] In one embodiment, network processor 1206 includes a PCI interface to a processor 1210. Processor 1210 may host certain applications, such as, for example, firewall applications. Processor 1210 may perform these functions with the use of hard disk 1211, DRAM 1213 and console port 1211. Console port 1211 may provide access to a monitor or keyboard or other peripheral device. In one embodiment, processor 1210 comprises a Pentium Processor manufactured by Intel Corporation of Santa Clara, California.

[00140] In one embodiment, network processor 1206 executes software instructions, which performs the 802.11 MAC layer. Network processor 1206 may also execute a wireless LAN configuration module to configure the wireless LAN network, a priority traffic administration (e.g., traffic shaping) module, a management software (e.g., Cisco IOS), a security protocol (e.g., 802.1x) module, and a VPN/firewall module. Processor 1210 executes a location tracking module to perform the location tracking. Processor 1210 may also execute one or more of the following software modules: clustering/HA, RADIUS/DHCP (remote authentication dial-In user service/dynamic host configuration protocol), session mobility, third party applications, XML (extensible markup language) Web services, user administration software, and network management software.

Reconfiguration of the Communication System

[00141] A technique described herein allows for the performance of an automatic site survey to reconfigure the wireless communication network. As part of the process, the repeaters in essence cause their own reconfiguration by providing information to the switch

that the switch uses to determine whether reconfiguration is necessary. In one embodiment, as a result of performing the reconfiguration process, one or more repeaters may change their state from activated, deactivated, or hot standby to another state and/or change their transmitter power level and/or receiver sensitivity. When in the activated state, a repeater is able to receive packets from sending devices (e.g., mobile devices in the network) and transmit packets to those devices. When in the deactivated state, a repeater is not able to receive packets from nor transmit packets to other devices (e.g., mobile devices in the network). When in the hot standby state, a repeater is able to receive packets from sending devices but not transmit packets to those devices. It is possible that a repeater may not change its state as part of the reconfiguration process, but may change its transmit power level and/or its receiver sensitivity.

[00142] Reconfiguration may also occur by having one or more repeaters change their channel numbers. The reconfiguration of the network includes turning on and off repeaters and adjusting transmitter power levels and receiver sensitivity. The reconfiguration occurs periodically. Reconfiguration may occur after a predetermined period of time (e.g., an hour) or a predetermined amount of activity. The reconfiguration may occur in response to an event. For example, if the activity of a repeater receives a predetermined number of packets within a predetermined period of time or the rate of packet reception increases by a predetermined amount, then the reconfiguration may be performed. As another example, the event may comprise a mobile station entering a particular location (e.g., a conference room) where a repeater is located and not on (thereby causing the system to be reconfigured to have the repeater activated). In one embodiment, when the event occurs, an alarm in the switch is triggered, causing the switch to run the reconfiguration process.

[00143] Figure 12B is a flow diagram of one embodiment of a process for reconfiguring the wireless communication system. Referring to Figure 12B, exemplary process 1250 begins by each repeater that is activated or in the hot standby state sending the received signal

strength indication, as set forth above, along with the SNR for each packet received cleanly to the switch (processing block 1251). As described above, the received signal strength and SNR may be determined on a packet-by-packet basis. Also as discussed above, communications between the repeater(s) and the switch(es) occur via a wired connection (e.g., an Ethernet connection) and/or may be through a level 2 network. Note that in another embodiment, such communications may be performed, at least in part, wirelessly using a different protocol than the protocol used between the mobile stations and the repeaters.

[00144] In response to sending the packet(s), the switch receives the packet(s) (processing block 1252) and determines the amount of wireless communication activity each repeater is experiencing (processing block 1253).

[00145] More specifically, the repeater receives a packet and embedded in the packet header is the Ethernet MAC address of the mobile station. When the repeater forwards that packet to the switch, it attaches the received signal strength and SNR values. In response to the packet, the switch is able to open up the packet and determine that the packet is from another unique IP address and, thus, another unique user. Based on this, the switch determines the density of unique users on a particular repeater. In other words, the switch determines the number of unique users (mobile stations) sending packets that are being received by an individual repeater. The switch may use a database to maintain this information. This database may be the location tracking database described above.

[00146] Based on the location and density of the repeaters as tracked by the switch, using the information sent from the repeater(s), the switch determines which repeaters to activate, deactivate, or move to the hot standby state (processing block 1254). The switch also determines the transmitter power levels for the repeaters that are activated (processing block 1255). The transmitter power levels are the power levels used by the repeaters when transmitting packets wirelessly to other devices in the network. The switch may also adjust the receive sensitivity of one or more of the repeaters (processing block 1256). The switch

may also adjust the channel numbers of one or more repeaters. In one embodiment, the switch causes these changes to be made by sending control commands to the repeater over, for example, a wired connection (e.g., the Ethernet connection).

[00147] Thus, if the switch determines that a particular repeater is to be activated (the repeater can receive and transmit), deactivated (the repeater cannot receive nor transmit), or placed in hot standby mode (the repeater can receive but cannot transmit), that changes to the repeater's transmitters power level and/or the repeater's receiver sensitivity, or that changes to the repeater's channel number are necessary, then the switch sends a command to the repeater specifying the desired action.

[00148] In one embodiment, if the number of unique users being received cleanly by a repeater in a hot standby state is above a threshold, then the switch activates the repeater.

[00149] This reconfiguration process has a number of advantages over the prior art. For example, as part of the reconfiguration process in the prior art, an access point may have to be moved. This is because there are typically no additional access points in the area that are not already being used because of their expense. In contrast, because repeaters are generally cheaper devices, many more of them may be distributed throughout the network, even though they are not going to be used all the time. Thus, when there is a need for additional capacity, one of the repeaters that is not currently activated can be activated.

[00150] In one embodiment, the reconfiguration of the wireless communication system may include changing the transmit power levels of the mobile stations. As with the reconfiguration described above, the purpose of this reconfiguration of the mobile station is to improve network capacity. The improvement to network capacity may be due to a reduced interference to repeaters and other mobile stations in adjacent coverage cells that a mobile station causes because its transmit power level is changed.

[00151] The reconfiguration of the mobile stations may occur in response to the switch examining the interference in a particular area and comparing this interference with a

predetermined amount of interference (e.g., a threshold). The predetermined amount of interference may be based on an allowable amount of interference for the wireless communication system or an allowable amount of variance from the allowable amount of interference.

[00152] The switch (or other control entity) determines the amount to change the transmit power level. In order to determine the amount of change to a particular transmit power level, the switch initially determines what the current transmit power level is. In one embodiment, the switch sends a query as a control message to the mobile station to obtain the transmit power level of the mobile station. Alternatively, the switch maintains a list (e.g., a database) of the transmit power levels of the mobile stations and accesses the list to obtain the transmit power level for a particular mobile station. The switch may obtain this information from the mobile stations. In addition, the switch might also send a command to the mobile station to modify its power level on a percentage basis. This would not require the knowledge of a specific power level. For example, in one embodiment, the mobile stations send a control message to the switch at boot-up indicating their transmit power levels.

[00153] Once the current transmit power level has been obtained, the switch determines the amount to change the transmit power level. This may be based on the received signal strength (e.g., RSSI) of the packets received by the repeater currently assigned to the mobile station. For example, if the received signal strength is very high, yet the mobile station is causing interference (e.g., its packets are being received by one or more other repeaters), the switch may cause the mobile station to reduce its transmit power level to a predetermined level or by a predetermined amount (e.g., a percentage of its current transmit power level) because the effect of such a reduction would not prevent its packets from being received by its assigned repeater.

[00154] The change in the transmit power level may be performed in a number of ways. For example, in one embodiment, the switch controls the transmit power level of the mobile

station(s). In such a case, the switch may send a command message to the mobile station, via a repeater, to cause the mobile station to adjust its transmit power level. The command could indicate that the mobile station should increase or decrease its transmit power level. Alternatively, such a command could come from a repeater or a control entity in the communication system other than the switch.

An Exemplary MAC Software Architecture

[00155] Figure 13 is one embodiment of a distributed MAC architecture. The 802.11 MAC layer is distributed between the switch and a number of the repeaters connected to the switch. On one side, the MAC is terminated on the switch and on the other side the MAC is terminated on the repeaters. Thus, in this way, the distributed architecture is “one to many” relationship.

[00156] In one embodiment, the MAC sublayer on the repeater is engaged in performing real time functions related to the time synchronization (BEACON, PROBE request/response processing), receiving and transmitting 802.11 frames, including acknowledgment of the received frames.

[00157] The MAC sublayer on the switch is centralized and controls multiple repeaters. In one embodiment, the MAC sublayer on the switch includes centralized management of the mobile stations and handles mobile stations in power save mode.

[00158] In one embodiment, the switch runs multiple instances of the MAC sublayer on the switch. In this manner, the switch may support multiple, separate logical groupings of repeaters on the switch. Each grouping may be based on channel frequency such that each group is associated with a particular frequency. The frequency need not be unique to all the frequencies of all the groupings (e.g., some groups use the same frequency and other groups do not use that frequency). The groupings may be created based on channel numbers.

[00159] By being able to run multiple instances of the MAC sublayer of the switch, the architecture offers flexibility when configuring the wireless communication system and individual embodiments that allows at least one of the following benefits. First, tuning of the size of the RF coverage per logical grouping of repeaters. Second, the roaming of the stations is easy to control. Third, the management of mobile stations in power save mode is centralized. That is, the frames for the mobile stations in power save mode are buffered in the MAC sublayer on the switch and can be exchanged between other instances of the MAC sublayer on the same switch (between MAC instances) when the mobile station in power save mode is roaming.

[00160] Referring to Figure 13, each of the units may be implemented in hardware, software, or a combination of both. Data_SAP unit 1301 exchanges messages with the LLC (logic link control) layer, conveying MSDUs (MAC service data units) from and to the LLC layer. Fragmentation unit 1302 performs fragmentation of outgoing MPDUs and MMPDUs (MAC management protocol data units). In one embodiment, since the sending of the fragmented PDU (protocol data unit) by a repeater has some timing constraints, the fragmented PDUs between the switch and the repeater are transferred in one tunneling protocol message. The tunneling protocol covers this case by putting a number of fragments in the tunneling protocol header. Power save unit 1303 performs power save device management, including TIM (Traffic Indication Map) management, in which TIM is sent to the repeaters periodically. The repeaters use the updated TIM for buffering of unicast MPDUs for mobile stations in power save mode. In one embodiment, the switch maintains buffered unicast PDUs for all mobile stations in power save mode. Broadcasts and multicast PDUs are not buffered at the switch and are sent to the repeaters to be sent out immediately after any beacon containing a TIM element with a DTIM (delivery traffic indication message) count field with a value of 0. Power save unit 1303 also performs PS-Poll request and response handling.

[00161] Routing unit 1305 routes data frames to MAC Data SAP (service access point) unit 1301 and management inbound frames to management_SAP unit 1309. De-fragmentation unit 1304 performs de-fragmentation of inbound frames. Management SAP unit 1309 includes an interface to MIB (management information base) unit 1308 and MLME (MAC sub-layer management entity) service unit 1307. MLME services unit 1307 handles the incoming associate and re-associate frames, as well as disassociate requests, and processes authentication and de-authenticate requests and generates authentication and de-authenticate response frames.

[00162] MIB management unit 1308 performs get and set functions to get and set parameters of the repeater and performs reset functions to reset all the parameters of a repeater and return the parameters to default values. The above processes may be performed using a tunneling protocol between a switch and the respective repeater.

[00163] With respect to block tunneling protocol layer 1306, both MPDUs and MMPDUs frames between the switch and the repeater are transferred by the tunneling protocol. In one embodiment, the 802.11 frames are encapsulated into Ethernet frames. In one embodiment, the tunneling protocol header is placed after the Ethernet header. This protocol transfers both data and management frames as well as special defined tunneling protocol control messages.

[00164] On the repeater, transmit unit 1311 transfers frames from MAC to PHY transmitter, generates FCS (frame check sequence), inserts timestamps in the beacons and probe responses, performs DCF timing (SIFS, DIFS, EIFS), handles ACK, RTS, CTS, and performs a back-off procedure.

[00165] Receive unit 1312 transfers frames from PHY to MAC, receives the MPDUs from the PHY, calculating and checking the FCS value (frames with valid FCS, length and protocol version are sent for receive filtering). Receive unit 1312 also filters valid received frames by destination address, and BSSID (basic service set identification) for group destination addresses, as well as handles ACK, CTS and RTS. Other functions include detection of

duplicate unicast frames, updating the NAV (network allocation vector) using Duration/ID value from 802.11 frames, maintenance of the channel state based on both physical and virtual carrier sense, time slot reference generation, and providing Busy, Idle & Slot signals to transmission.

[00166] Synchronization unit 1313 processes the MLME start request in which it starts a new BSS (basic service set) and set all parameters for a beacon frame. Synchronization unit 1313 generates beacon frames periodically and handles Probe request and response frames.

[00167] Repeater management unit 1314 relays all MIB set/get requests, start requests, reset requests, request/confirm characteristic commands to a proper block on the repeater.

[00168] With respect to block tunneling protocol layer 1310, frames for both MPDUs and MMPDUs between the switch and repeater are transferred by the tunneling protocol. The frames are encapsulated into the Ethernet frames and the tunneling protocol header is placed after the Ethernet header. This protocol transfers both data and management frames as well as special defined tunneling protocol control messages.

An Exemplary Switch Software Architecture

[00169] The switch contains the switching and management planes. Figure 14 illustrates one embodiment of the switching plane. Referring to Figure 14, the switching plane 1400 contains the switch MAC sublayer 1402 (i.e., the upper MAC), a switch management entity (SwME) 1401 and a switching layer 1403. The switching layer 1403 interfaces with the Ethernet drivers 1404 and performs the switching function. The Ethernet drivers 1404 are connected to the 10/100 BT ports of the switch (PORT1 to PORT24) or connected to another Ethernet switch with its uplink connected to the Gigabit interface 1406 on the switch. The simulator 1405 may also be connected to any of these ports. In one embodiment, in order to support this kind of abstraction, the tunneling protocol header contains the number of the Ethernet port assigned for use with the repeater.

Data Traffic Procedures

[00170] Figures 15-18 illustrate the communication network and exemplary data traffic process. Referring to Figure 15, switch 1501 is shown coupled to router 1502 and repeaters 1-3, via ports 1-3. Stations (STA) 1-4 are mobile stations that communicate wirelessly with the repeaters 1-3. Router 1502 is also shown coupled to computer system 1503.

[00171] Figure 16 illustrates an exemplary process for transferring data traffic from a mobile station to a desktop computer system. Referring to Figure 16, repeater 1602 receives the one or more 802.11 data frames (packets) and encapsulates each received 802.11 data frame into one or more Ethernet packets, adding an Ethernet frame header and a tunneling protocol header to each Ethernet packet. Thereafter, repeater 1602 sends the Ethernet frames (packets) to switch MAC sublayer 1603 on the switch. At the switch, switch MAC sublayer 1603 processes the Ethernet frames by stripping off the 802.11 MAC header and tunneling protocol headers and switches Ethernet frames (packets) with encapsulated IP packets to the proper switch port. Switch MAC sublayer 1603 sends the Ethernet frames (packets) to router 1604 (backbone). Router 1604 routes each Ethernet frame to a destination, such as, for example, computer system 1605.

[00172] Figure 17 illustrates an exemplary process for transferring data traffic between two mobile stations. In this case, the destination address is another mobile station address and the switch MAC sublayer processes both the 802.11 and tunneling protocol headers and switches the packet to the proper port. Referring to Figure 17, a first station, station 1701, sends 802.11 data frames to a first repeater, repeater 1702. Repeater 1702 receives the 802.11 data frame and encapsulates the 802.11 frames into Ethernet frames, including adding an Ethernet frame header and tunneling protocol header to each 802.11 frame. Repeater 1702 sends the encapsulated 802.11 data frames to switch MAC sublayer 1703. Switch MAC sublayer 1703 processes the 802.11 data frames and tunneling headers and switches Ethernet frames to the

repeater (repeater 1704 in this example) handling the destination station (station 1705 in this example). Switch MAC sublayer 1703 encapsulates the 802.11 data frames into Ethernet frames and sends them to repeater 1704. Repeater 1704 receives the encapsulated 802.11 data frames and sends the 802.11 data frames to station 1705.

[00173] Figure 18 illustrates an exemplary process for transferring data traffic from a desktop computer system to a mobile station. Referring to Figure 18, computer system 1806 encapsulates IP packets into Ethernet frames. For the first IP packet destined to a mobile station, the router starts an ARP (address resolution protocol) procedure in order to obtain the corresponding MAC address. Router 1805 sends an ARP request to switch MAC sublayer 1804 to request the MAC for this IP broadcast. Switch MAC sublayer 1804 encapsulates the ARP request into an 802.11 packet and then encapsulates this packet into an Ethernet packet, essentially creating a new Ethernet frame with an embedded 802.11 MAC header and tunneling protocol header. Switch MAC sublayer 1804 broadcasts this packet to all repeaters, repeaters 1802-1803 in this example, which then rebroadcast it for the desired mobile station to receive. The mobile station, station 1801, with the IP address contained in the ARP request sends an ARP response with its MAC address. Repeater 1802 receives the ARP response and encapsulates the 802.11 frames into Ethernet frames, adding an Ethernet frame header and tunneling protocol header. Repeater 1802 sends the encapsulated ARP response to switch MAC sublayer 1804, which strips off the 802.11 MAC header and switches the Ethernet frame with encapsulated ARP response packet to the backbone port.

[00174] After this procedure, the router takes the station MAC address from the ARP response and routes all IP packets for this mobile station as described above. Since the switch MAC sublayer has the configuration information about MAC and IP addresses, the ARP response could come from the switch.

Management Procedures

[00175] There are a number of management procedures supported by the distributed MAC architecture. In one embodiment, these include starting up the switch, resetting the MAC, starting a new BSS, synchronization, authentication, and de-authentication, association, disassociation and re-association.

[00176] With respect to starting up the switch, the switch is started by the switch management entity (SwME). To configure and start the switch and the repeaters, the SwME issues commands to the switch MAC sublayer on the switch. The commands intended for the repeaters are transferred using the tunneling protocol. Layers of the tunneling protocol are running on the switch and the repeaters.

[00177] With respect to MAC reset, the switch and repeaters cooperate to perform a reset of the MAC. Since the MAC is distributed between the switch and repeaters, the reset process is modified to support this architecture. In one embodiment, the switch management entity sends a reset request to each of the repeaters as part of a tunneling protocol process and receives a reset response indicating if the reset was successful. The reset process may set the MAC to initial conditions, clearing all internal variables to the default values. MIB (management information base) attributes may be reset to their implementation-dependent default values.

[00178] With respect to the start process, the switch management entity requests that the MAC entity start a new BSS. The switch management entity generates the request to start an infrastructure BSS (basic service set) with the MAC entity acting as an access point and sends it to all MAC entities where the switch is acting as a multiple access point. Each repeater responds with an indication as to whether the start process was successful.

[00179] With respect to synchronization, the synchronization process determines the characteristics of the available BSSs and allows for synchronizing the timing of a mobile station with a specified BSS (switch MAC entity). In one embodiment, the synchronization process begins with an instance of the switch MAC sublayer generating a beacon frame,

which is encapsulated and sent to the repeaters periodically. The repeater updates the timestamp of the beacon frame before sending the beacon frame in the air. Based on the beacon frame, the mobile station synchronizes its timers.

[00180] The switch management entity also causes authentication to establish a relationship between a station MAC sublayer and the instances of the switch MAC sublayers. In one embodiment, a mobile station is authenticated if its MAC address is in the access list on the switch. Similarly, de-authentication is supported to invalidate an authentication relationship with a switch MAC entity. In one embodiment, de-authentication is initiated by the mobile station. In this case, the instance of the switch MAC sublayer on the switch associated with the repeater assigned to the mobile station updates the station state as maintained by the switch. The result of de-authentication is that the state of the mobile station is listed in the switch as unauthenticated and unassociated.

Association

[00181] Data frames for a mobile station are forwarded from the repeater that has the token for the mobile station. In one embodiment, if a repeater without the token receives the data frames, it forwards only a short frame with the RSSI (in the tunneling protocol header) to the switch. The switch keeps track of the RSSI for the mobile station. If the repeater without the token has better reception and if the repeater with the token has “high” error rate, the switch may re-assign the token. The RSSI and token are part of the tunneling protocol header. The token assignment occurs within the association process.

[00182] Figure 19 is a data flow diagram of one embodiment of an association and token assignment process. Referring to Figure 19, an association request is generated by a mobile station and sent by the mobile station, via the mobile station MAC. Repeater 2 has the token for the mobile station. Therefore, repeater 2 encapsulates the association request, along with its RSSI and BSSID, into an Ethernet packet and sends the encapsulated packet to the switch.

Repeater 1, which does not have the token for the mobile station, forwards a short frame with the RSSI in the tunneling protocol header.

[00183] The switch takes the RSSIs for the two identical frames and determines which one is stronger. Based on which is stronger, the switch either allows the repeater that has the token and station MAC for the mobile station to keep them (e.g., repeater 2) or reassigns them to the repeater with the higher RSSI (e.g., repeater 1). In either case, the switch sends an association response encapsulated in an Ethernet packet with the token and association ID to the repeater, which de-encapsulates it and forwards it to the mobile station, via the mobile station MAC.

Re-association

[00184] The following exemplary procedure describes how a mobile station becomes re-associated with another switch MAC entity (logical access point). Figure 20 is a block diagram of two MAC sublayer instances in a switch. Referring to Figure 20, two (or more) instances of the switch MAC sublayer run on the switch (offering the access points (APs) inside the same switch). Each instance has its own BSSID (basic service set identification) (e.g., the MAC address of the MAC instance). Both MAC instances are managed by the same switch management entity (SwME). The SwME manages these as multiple access points (APs) inside the switch. In one embodiment, communication between MAC instances is through the SwME. Both MAC instances as well as the switch management entity (SwME) reside on the same switch. Communication between the MAC instances can be direct or through the SwME. In one embodiment, the SwME has knowledge of all MAC instances and is involved in this communication. Thus, the switch acts as a distribution system containing multiple switch MAC sublayer instances (multiple logical access points) in which roaming is centralized in the switch.

[00185] In one embodiment, the association request from the mobile station is encapsulated and sent by the repeater to the switch. The association request with the BSSID of the first MAC sublayer instance is sent from the second MAC sublayer instance through the SwME to the first MAC sublayer instance. As a result, the first MAC sublayer instance generates a response representing that mobile station has been already associated with the first MAC sublayer instance. Using this process, the station does not have to go again through authentication procedure and it can be automatically associated with the second MAC sublayer instance. When the second MAC sublayer instance receives the response, the station becomes associated with the second MAC sublayer. Thus, when the station roamed, the handover procedure is performed in the switch. Therefore, the switch acts as a complete distribution system with multiple logical access points.

[00186] As described above, when a station roams between two MAC sublayer instances (logical access points) inside one distribution system, there is only one repeater controlled by one MAC sublayer instance. In one embodiment, a mobile station can roam from one repeater to another repeater controlled by the same MAC sublayer instance (logical access point) without a need to associate again, and only the token re-assignment procedure described herein has to be performed. In one embodiment, the station is not aware of the token re-assignment procedure.

[00187] If a mobile station moves from one repeater belonging to one logical access point (one MAC sublayer instance) to a second repeater belonging to a second logical access point (second MAC sublayer instance), the station has to be re-associated and the token re-assignment procedure has to be performed. The handover procedure is performed in the switch. Again, the station is not aware of any token assignment procedures.

[00188] Note that mobile stations are associated with switch MAC sublayers instances not with a repeater. If a station is controlled by a repeater, the repeater has a token for that station. All repeaters controlled by a particular MAC sublayer instance are associated with a station if

the station is associated with that MAC sublayer instance, and only one repeater has a token for that station.

[00189] A user can configure the switch to have any number of MAC instances. In one embodiment, this is configured using a parameter. Also configurable is which repeater belongs to MAC instance. For example, if the switch has 64 ports, it can be configured to act as 8 access points (8 upper MAC instances running concurrently), with 8 repeaters per access point (one upper MAC sublayer controlling 8 repeaters).

[00190] Figure 21 is a data flow diagram of one embodiment of a re-association process. Referring to Figure 21, a mobile station SME (station management entity) generates a re-association request and sends it to a repeater, repeater 4 in this case, along with its BSSID via the mobile station MAC. In one embodiment, the mobile station knows that it needs to make a re-association request because it has received a BEACON frame with different BSSID (i.e., a different MAC instance), indicating that the mobile station had been roaming. The repeater receives the re-association request, encapsulates the packets of the re-association request with the RSSI into an Ethernet packet, and sends the Ethernet packet to the instance of the switch MAC sublayer associated with the repeater. In response thereto, the instance of the switch MAC sublayer generates an indication to the switch management entity indicating that a re-association request has been made.

[00191] In response to the indication, the switch management entity causes a new AID (association id) to be assigned to the mobile station, a token for the mobile station to be assigned to a new repeater, and the previous token assignment to be deleted. In one embodiment, the association identifier (AID) is a number (value between 0 and 2007) assigned to a mobile station by the switch or an access point during the association procedure. It is an 802.11 standard defined parameter. After the station is associated, the station inserts the AID in every message. More specifically, the switch management entity updates the entry for the mobile station in the access list, including setting the new access point address to the

address of the instance of the switch MAC sublayer associated with the repeater. The switch management entity also assigns a token and an association ID.

[00192] The switch management entity sends a delete token command to the instance of the switch MAC layer associated with the repeater previously assigned to the mobile station, which the instance of the switch MAC layer forwards to the repeater (repeater 3 in this case).

[00193] The instance of the switch MAC sublayer (upper MAC 2 in this case) associated with the repeater that forwarded the re-associate request (repeater 4 in this case) sends a re-associate response frame to the repeater with the token, association ID, and an indication that the re-association was successful. The repeater de-encapsulates the packet, stores the mobile station MAC token, and forwards the de-encapsulated re-associate response frame to the mobile station MAC with the association ID and the successful indication.

Disassociation

[00194] A mobile station may request disassociation with a specified peer MAC entity that is acting as a logical access point. The mobile station may request this due to inactivity or because a switch is unable to handle all currently associated mobile stations, etc.

Figure 22 is a flow diagram on one embodiment of a disassociation process. Referring to Figure 22, a disassociation request is generated by the SME (station management entity) on the mobile station and sent by the mobile station MAC as a disassociate request frame with the BSSID (i.e., the instance identifiers). The BSSID is a basic service set identifier representing the MAC address of an upper MAC instance. Each repeaters that receives the disassociate request frame without errors encapsulates it with its RSSI and forwards it to the switch, regardless of whether it has the token for the mobile station. In response to the receiving the disassociate request frame, the switch MAC sublayer determines whether the mobile station is in the access list and changes the state of the mobile station in the access list to authenticated and unassociated, removes all parameters from the access list entry for the

mobile station, and deletes the token and association ID. In one embodiment, the access list is a dynamically created hash table containing records for all authenticated stations, in which each record contains a station MAC address, association identifier, BSSID, a station state, and a repeater port number which has station token. In other words, on the switch MAC sublayer, the state of the mobile station is updated and its AID is deleted. The switch then sends a disassociate response frame encapsulated in an Ethernet frame to the repeater having the token. Embedded in the tunneling protocol header of the frame is a tunneling protocol command to delete the token, which causes the repeater having the token to delete the token. Thereafter, the repeater that deleted the token sends the de-encapsulated disassociate response frame to the MAC of the mobile station with an indication that disassociation was successful.

[00195] In one embodiment, this process can be initiated by the switch management entity. This can happen if the switch decides to disassociate the mobile station because of inactivity or because a switch is unable to handle all currently associated mobile stations.

An Exemplary Communication Network

[00196] Figure 23A is a block diagram of an embodiment of a communication system. The wireless components of exemplary system 2300 may be operating at a substantially the same frequency, which will be described in details further below. In one embodiment, system 2300 includes a switch 2301 having one or more switch ports which may be coupled to one or more repeaters, such as repeaters 2302 to 2305. Repeaters 2302 to 2305 may be wireless communication devices capable of receiving and transmitting data wirelessly with one or more mobile stations, also referred to as packet antennas. Repeaters 2302 to 2305 may wirelessly communicate with one or more mobile stations, such as mobile stations 2314 to 2316. A mobile station may be able to wirelessly communicate with more than one repeater. For example, mobile station 2316 may be able to communicate with repeaters 2302 and 2304. Each of the repeaters 2302 and 2304 may include a received signal strength indicator (RSSI)

that allows switch 2301 to determine which repeater is the closest repeater with respect to mobile station 2316. In one embodiment, the repeaters and the mobile stations are operating at substantially the same frequency. In this embodiment, switch 2301 is responsible for managing transmission of packets to the mobile stations to avoid data collisions.

[00197] In one embodiment, each of the repeaters 2302 - 2305 receives one or more packets of wirelessly transmitted packets from the mobile stations. In one embodiment, each of the repeaters forwards to the switch each packet of the wirelessly transmitted packets that each repeater had received at a received signal strength specified by the switch, such as, for example, a threshold set by the switch. Alternatively, only the repeater with the token for a mobile station forwards the packets from the mobile station to the switch. The repeaters may forward to the switch RSSI values for all packets received without error. In one embodiment, the repeaters that are not the primary repeater corresponding to the mobile station (e.g., those without the token for the mobile station) only send the respective RSSI through a particular RSSI information message using tunneling protocols, such as, for example, exemplary RSSI information message 4200 shown in Figure 42, without data payload, which is carried by the primary repeater. The primary may transmit its RSSI to the switch using exemplary RSSI information message 4200 when there is no data payload received from the mobile station. However, when the primary repeater receives data payload (e.g., 802.11 data, control, or management message) from the mobile station, it may encapsulate its RSSI within a message that carries the data payload, such as, for example, exemplary messages shown in Figures 39A-39C.

[00198] In one embodiment, the repeaters are grouped and the switch handles each group of repeaters separately. Even so, if a mobile station moves to a location in which a different repeater in a different group is associated with the mobile station, any data buffered by the switch may be forwarded to the mobile device through the new repeater using a single data transfer within the switch.

[00199] In addition to wireless components, according to one embodiment, a repeater may further include one or more ports to allow another station to couple to the switch. For example, repeater 2303 may include a wired port to connect with an IP ready phone which may be able to communicate with other phone systems using voice over IP (VoIP) techniques, while repeater 2303 is able to communicate wirelessly with other mobile stations, such as mobile station 2315. Similarly, repeater 2305 may include a wired port to connect a wired station, such as a conventional desktop station 2312.

[00200] Compared to traditional 802.11 access points, in one embodiment, the repeaters are technologically light, simply providing a portal from the RF air medium into a wired network where network convergence occurs. According to one embodiment, the repeaters do not perform any or very little packet processing themselves, the repeaters merely pass data, management and control frames back and forth between wireless clients and switch 2301. Nor do the repeaters perform any of the access control, security, or management functions of the conventional 802.11 access points. Instead, switch 2301 performs most of these functionalities, which will be described in details further below.

[00201] Furthermore, according to another embodiment, switch 2301 may include one or more ports suitable to connect a local area network (LAN), such as a LAN with one or more standard wired desktop systems 2311, to switch 2301. According to yet another embodiment, the local area network coupled to switch 2301 may be another wireless network using, for example, a third party access point device 2306.

[00202] According to one embodiment, an uplink of switch 2301 may be coupled to an enterprise backbone, such as enterprise backbone 2308, through a router. Enterprise backbone 2308 may include one or more local area networks, such as Intranet 2309 having one or more servers and clients 2320. A network management console 2307 may be coupled to the enterprise backbone 2308 to access switch 2301. For example, a network administrator may use network management console 2307 (e.g., via a graphical user interface (GUI) or a

command line interface (CLI)) (e.g., telnet) to access and to configure switch 2301 on behalf of one or more client nodes (e.g., mobile stations or wired stations), including Ethernet routing information and security policies associated with each of the client nodes. In one embodiment, the Ethernet routing information and security policies may be stored in a table within switch 2301, which will be described in details further below.

[00203] According to one embodiment, switch 2301 includes one or more ports 2321 (also referred to as public interfaces or public ports) coupled to a separate network 2309, which will be described in details further below. Separate network 2309 may be a publicly accessible network, such as a wide area network (WAN) (e.g., Internet) or a publicly accessible dedicated network. Separate network 2309 may be a private network, such as a corporate Intranet network. When switch 2301 receives a packet from one of the client nodes, either a wired node or a mobile station, switch 2301 may route the packet either to publicly accessible network 2309 via one or more public ports 2321 or to enterprise backbone 2308 via uplink 2322, based on the Ethernet routing information and the security policies associated with the client nodes.

[00204] According to one embodiment, the repeaters (e.g., repeaters 2302-2305) communicate with switch 2301 according to a tunneling protocol within an Ethernet protocol, such as Ethernet 802.3 protocol, and the repeaters communicate with the mobile devices via a wireless communication protocol, such as an IEEE 802.11 protocol.

An Exemplary Wireless Communication System Using A Tunneling Protocol

[00205] Figure 23B is a block diagram illustrating an exemplary embodiment of a wireless communication system using a tunneling protocol. In one embodiment, similar to exemplary system 2300 of Figure 23A, exemplary system 2350 includes, but not limited to, switch 2351 communicatively coupled to one or more repeaters 2352 and 2353 via a wired or wireless communication media 2357, such as an Ethernet 802.3 communication media. Repeaters

2352 and 2353 communicate with one or more mobile stations 2354-2356 via a wireless communication media 2358, such as an IEEE 802.11 communication media.

[00206] According to one embodiment, communications between switch 2351 and repeaters 2352 and 2353 are carried out via a tunneling protocol, also referred to as ATP which is encapsulated within an Ethernet packet (e.g., an Ethernet 802.3 packet or a wireless packet, such as, for example, an IEEE 802.16 packet). For example, communications between switch 2351 and repeater 2352 are performed via a logical tunnel 2359 within the Ethernet communication protocol. When switch 2351 needs to transmit a packet (e.g., a data or control packet) to a repeater, such as repeater 2352 or a mobile station, such as mobile station 2354, switch 2351 converts the packet into a tunneling protocol packet, also referred to as an ATP packet 2361, conforming to a tunneling protocol, such as ATP protocol. Switch 2351 then encapsulates the ATP packet 2361 within an Ethernet packet 2360 and transmits the Ethernet packet 2360 to repeater 2352. If the destination of ATP packet 2361 is repeater 2352, repeater 2352 extracts ATP packet 2361 from Ethernet packet 2360 and performs operations according to the extracted ATP packet 2361.

[00207] If it is determined that the destination of ATP packet 2361 is a mobile station, such as mobile station 2354, based on, for example, the MAC address of the mobile station, repeater 2352 extracts ATP packet 2361 from Ethernet packet 2360. Repeater 2352 then extracts the ATP packet from Ethernet packet 2360 to recover the corresponding 802.11 packet. The 802.11 packet is then transmitted to mobile station 2354 via the 802.11 communication media.

An Exemplary Tunneling Protocol Format

[00208] The tunneling protocol frames, also referred to as airflow tunneling protocol (ATP) frames, are designed to tunnel through a standard Ethernet or 802.3 frame. This allows the frames to be routed through a layer 2 subnet from source to destination, using standard

Ethernet/802.3 switching components. As a tunneling protocol, ATP frames are designed to co-exist with Ethernet/802.3 frames on the same network segment. Thus, wireless and wired traffic can share the same connection between a repeater (also referred to as a packet antenna) and a switch.

[00209] In one embodiment, ATP frames with a high priority may be given precedence over other request frames made at the same end station (e.g., a switch, a repeater, or a mobile station) or at other stations attached to the same LAN. The resultant transmission delay may include a queuing delay until the frame becomes first in line for transmission on the LAN by a repeater or a mobile station and an access delay for transmission of the frame.

[00210] In one embodiment, there are two formats used in LANs when representing MAC address values in MAC user data: canonical format and non-canonical format. The native format used in MAC addresses where the bit transmission order is least significant bit first is the canonical format. In the canonical format, the least significant bit of each octet of the standard hexadecimal representation of the address (see IEEE Standard 802-1990) represents the least-significant bit (LSB) of the corresponding octet of the canonical format of the address. In the non-canonical format, the most significant bit (MSB) of each octet of the standard hexadecimal representation of the address (see IEEE Standard 802-1990) represents the least-significant bit of the corresponding octet of the canonical format of the address.

[00211] Figure 24A shows a typical Ethernet frame format. In one embodiment, all ATP Frames are transmitted as Ethernet/802.3 frames at layer 2. These frames also conform to the IEEE 802.1Q VLAN standard. Referring to Figure 24A, the Ethernet frame includes a type field 2401 and data field 2402. Type field 2401 is used to identify the type of the data field 2402. Since the ATP leverages the 802.1Q VLAN standard, according to one embodiment, every ATP frame includes the value of '0x8100' in its type field, such as type field 2401. The selection of this IEEE standard is based on at least one of the following benefits:

- VLANs are supported over all IEEE 802 LAN MAC protocols, and over shared media LANs as well as point-to-point LANs.
- VLANs facilitate easy administration of logical groups of stations that can communicate as if they were on the same LAN. They also facilitate easier administration of moves, adds, and changes in members of these groups.
- Traffic between VLANs is restricted. Bridges forward unicast, multicast, and broadcast traffic only on LAN segments that serve the VLAN to which the traffic belongs.
- As far as possible, VLANs maintain compatibility with existing bridges and end stations.
- If all ports that provide bridge services for other networks (e.g., other LANs) are configured to transmit and receive untagged frames, bridges will work in plug-and-play ISO/IEC 15802-3 mode. End stations (e.g., repeaters or mobile stations) will be able to communicate throughout the bridged LAN. This allows non-repeater Ethernet devices to connect directly to the switch, such as desktop computers, legacy Access Points, etc.
- 802.1p prioritization is included in the 802.1Q standard. This allows taking advantage of the Class Of Service provisions in the protocol to support bandwidth policing and monitoring.
- fully integrated support in Ethernet switching chips, including Broadcom's StrataSwitch II, BCM5616 Multi-Layer Switch chip.

[00212] Figure 24B is a diagram of an exemplary embodiment of Ethernet frame that is used to transmit an ATP frame. Referring to Figure 24B, in one embodiment, the exemplary frame includes VLAN tag field 2403, ATP header 2404, and ATP frame body 2405. The value of '0x8100' in the location of the type field (e.g., VLAN tag field 2403) is also called the tag protocol identifier (TPID) of a VLAN frame. In one embodiment, tagging of VLAN frames has at least one of the following purposes:

- To allow user priority information to be added to frames carried on IEEE 802 LAN MAC types that have no inherent ability to signal priority information at the MAC protocol level;
- To allow a frame to carry a VID (VLAN Identifier);

- To allow the frame to indicate the format of MAC Address information carried in MAC user data;

In one embodiment, since multiple VIDs are supported per port on an end station, traffic between them can be classified based on the VID value that is used.

[00213] Figure 24C is a diagram of an exemplary embodiment of a VLAN tag field, also referred to as tag control information (TCI) field. The exemplary VLAN tag field may be used as VLAN tag field 2403 of Figure 24B. In one embodiment, exemplary VLAN tag field 2403 includes a user priority field 2406, a CFI (canonical format indicator) field 2407, and a VLAN ID (VID) field 2408. It will be appreciated that the bit and field orders are not limited to the one shown, other specific formats may be implemented.

[00214] According to one embodiment, user priority field 2406 defines up to 8 priorities in accordance with IEEE Std 802.1p. Since the MAC layer service does not permit the reordering of frames with a given user priority for a given combination of destination address and source address, this field facilitates expedited traffic flow through the layer 2 network connecting a repeater and a switch.

[00215] According to one embodiment, CFI 2407 is a single bit flag that specifies the format of any MAC addresses that may be present in the ATP Frame Body. This bit may be set to zero to indicate that the canonical format is used. In one embodiment, for both canonical and non-canonical formats, the octet ordering is identical in both cases (e.g., the first octet is the leftmost octet of the address when written down using hexadecimal notation as defined in Section 5 of IEEE Std 802). However, the bit ordering within each octet may be different. For example, in the canonical format, the LSB of each octet of the standard hexadecimal representation of the address (see IEEE Standard 802-1990) represents the LSB of the corresponding octet of the canonical format of the address. The VLAN Identifier (VID) uniquely identifies the VLAN to which the frame belongs.

[00216] According to one embodiment, several VIDs are used for normal operations. These VIDs may be unique for each deployment. They may be assigned by IT professional during the initial system setup. The following is an exemplary list of VIDs that may be used in one embodiment by the ATP.

Description	VID
Untagged Desktop Traffic	0x06
Packets between one switch and another switch	0x01
Packets between one repeater and another repeater	0x02
Management and Control Packets between switch and repeater	0x03
Authorized Data Packets for normal routing	0x04
Unsecured Data Packets for routing untrusted packets	0x05

The VID for each type of transactions may be the same or different according to different embodiments. It will be appreciated that other values may be used.

[00217] Figure 24D is a diagram of an exemplary embodiment of an ATP header. The ATP header is used to send ATP messages to a destination endpoint. In one embodiment, exemplary ATP header 2404 includes, but not limited to, type field 2409, length field 2410, ATP type field 2411, ATP version field 2412, ATP message ID 2413, RSSI/Power field 2414, and transaction ID field 2415. It will be appreciated that the bit and field orders are not limited to the one shown; other specific formats may be implemented.

[00218] Referring to Figure 24D, according to the 802.1Q specification, type field 2409 may be implemented to contain either the length or the protocol type value for the remainder of the frame. In a particular embodiment, the Ethernet type (e.g., 0x1111) is used in type field 2409.

[00219] Length field 2410 is the ATP payload length and indicates the number of data bytes following the ATP header. To meet the Ethernet frame minimum 64-byte frame size requirement, the size of the frame following length field 2410 should not be less than 44. Zero-padded data may be used to maintain this requirement. This field does not include any padded data.

[00220] ATP type field 2411 specifies the type of ATP protocol for this frame. The ATP protocol can encapsulate a wide variety of protocol formats. Using ATP type field 2411, according to one embodiment, a unique tunneling protocol type representing the type of ATP protocol is identified. The tunneling protocol type may be an identification assigned by a related industrial group, such as IANA (internet assigned numbers authority). ATP version field 2412 indicates the version of the ATP protocol for this corresponding ATP tunnel type. RSSI/Power field 2414 is used to specify current RSSI or power level of an end node, when it is needed. For some Message IDs (specified in ATP message ID field 2413), the frame format includes a transaction identifier (TID) 2415. For non-ACK messages that include this field, an acknowledgment frame, such as, for example, exemplary acknowledgment message 3800 shown in Figure 38, is a response from the destination node. In one embodiment, the TID is a 6-bit value of which, the upper two bits may be used for data fragmentations to indicate whether a fragment is a first [0:0], a continuation [0:1], or a last [1:0] fragment.

According to one embodiment, the assigned value may be copied into the respective acknowledgment frame that is generated as a response.

[00221] ATP message ID field 2413 specifies a particular ATP message that is transmitted via a tunneling protocol within the Ethernet protocol. Figure 25A is a diagram of an exemplary list representing at least a portion of the ATP messages according to one embodiment. In one embodiment, exemplary list 2500 includes certain categories of messages, including, but not limited to, switch-to-repeater 2501, repeater-to-switch 2502,

switch-to-mobile-station 2503, mobile-station-to-switch 2504, switch-to-switch 2505, and repeater-to-repeater 2506. It will be appreciated that ATP messages are not limited to those shown in Figure 25A, more or less ATP messages may be implemented dependent upon the specific application.

[00222] Figure 25B is a diagram of an exemplary embodiment of an ATP frame, such as, for example, an initialization message for a repeater (e.g., message 2507 of Figure 25A having an ID of 0x0), which may be used to initialize a repeater. The reference numbers are maintained the same with respect to Figures 24B-24D for illustration purposes only. Referring to Figure 25B, exemplary ATP frame 2507 includes VLAN tag field 2403, ATP header 2404, and ATP frame body 2405, similar to those shown in Figures 24B-24D. The first two bytes of frame payload 2405 may be the parameter count. The remaining of the payload may include data parameters. Data parameters (including firmware) are carried in the frame payload in the form of 3-tuples.

[00223] Figure 25C is a block diagram of an exemplary embodiment of a packet in a 3-tuple format. In one embodiment, exemplary format consists of an identifier 2520, length (length of data following, not including identifier and length) 2521, and data 2522. In one embodiment, the data field is in big-endian format. Figures 26A and 26B show at least a portion of exemplary data that may be used as data identifier 2520, data length 2521, and data 2522.

Exemplary Discovery Processes of Repeaters

[00224] Figure 27A is flow diagram of an exemplary embodiment of a discovery process. Exemplary process 2700 may be performed by one or more processing logic that may include hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. Referring to Figure 27A, when repeater 2702 is activated (e.g., powered up and plugged into a port of switch 2701),

repeater 2702 sends a broadcast message to a network that includes switch 2701 and one or more other repeaters to announce that repeater 2702 is entering the network (operation 2703). In one embodiment, the broadcast message is a part of a heartbeat message regularly or periodically transmitted from repeater 2702 when it is activated. An exemplary heartbeat message frame used by a repeater is shown in Figure 34A. However, the format of the message is not limited to the one shown in the figure. It will be appreciated that other formats may be utilized.

[00225] Referring to Figures 27A and 34A, initially, according to one embodiment, heartbeat payload field 3402 may include a message (e.g., tunneled message) carrying an operating state of repeater 2702 to indicate that repeater 2702 is entering the network. For example, in a particular embodiment, heartbeat payload field 3402 includes an operating state message, such as, for example, message having an ID of 0x21 shown in Figure 26A to specify a current operating state of repeater 2702. An exemplary embodiment of operating states is shown in Figure 35. The operating states are not limited to those shown in Figure 35. More or less operating states may be implemented. In one embodiment, repeater 2702 may indicate it is operating in a discovery state (see, Figure 35). Alternatively, VLAN tag field 3401 may be left with a predetermined value or not used at all, indicating that repeater 2702 is entering the network. Furthermore, other fields, such as, for example, the transaction ID field, may be used for such purposes.

[00226] Once switch 2701 receives the initial broadcast message from repeater 2702, switch 2701 responds the message by downloading VLAN configuration information to repeater 2702 (operation 2704), which may be used by repeater 2702 for further communications in the network. In one embodiment, the response from switch 2701 is carried by a set data message (e.g., set data value message having an ID of 0x06 shown in Figure 25A) for a repeater. An exemplary set data message for a repeater is shown in Figure 36A. However, the format of the message is not limited to the one shown in the figure. It

will be appreciated that other formats may be utilized. The frame payload of the set data message may include data tuples that contains the VLAN configuration information (e.g., message having an ID of 0x1F shown in Figure 26A). An exemplary embodiment of VLAN configuration information is shown in Figure 37. Exemplary VLAN configuration 3700 includes a VLAN ID for each type of transactions 3701-3706. It will be appreciated that other types of transactions may be included.

[00227] Once repeater 2702 receives the VLAN configuration information from switch 2701, it stores the VLAN configuration information in a local storage and sends a response message with respect to the initialization message received from the switch. An exemplary embodiment of a response message is shown in Figure 36B. The payload of the response message may include the status to the initialization command and the current operating state of the repeater, such as, for example, a “not loaded” state as shown in Figure 35. In one embodiment, a logical value of 1 for the initialization status indicates success. Other values other than 1 indicate errors, which may be further described through optional error codes. Thereafter, repeater 2702 may change its heartbeat message to conform to one of the VLAN ID provided by the VLAN configuration information received from switch 2701 (e.g., VLAN ID 3401 of exemplary heartbeat message 3400 of Figure 34A may be changed to conform to one of the VLAN ID 3701-3706 of Figure 37) and transmits the new heartbeat message to the network (operation 2705).

[00228] Once switch 2701 receives the response with respect to the VLAN configuration information indicating a success, switch 2701 may start to download operating software (that may be used to program firmware) to repeater 2702 to establish an operating environment (e.g., an operating system) of repeater 2702 (operation 2706). In one embodiment, the software is downloaded to repeater 2702 via a set data message, similar to an exemplary set data message 3600 shown in Figure 36A. The software being downloaded is carried within the frame payload of the message. In one embodiment, the software may be divided in

multiple segments and the multiple segments may be sequentially downloaded (e.g., via multiple set data messages) to repeater 2702. For each segment received, repeater 2702 may responds with a response message, similar to exemplary set data response message 3650 shown in Figure 36B, indicating whether the data has been successfully received. The response message further includes an error code field that specifies a type of errors for any possible failure. In one embodiment, each set data message and its response message includes a sequence number to indicate the current sequence of the respective segment. If repeater 2702 determines that one of the segment or sequence has not arrived, repeater 2702 may return an failure status optionally with an error code via the response message similar to the one shown in Figure 36B. In this case, switch 2701 may retry to send the same segment again. Alternatively, repeater 2702 may abort the download and discard the software that has been downloaded, and resets it self (e.g., restart over again including discovery process, etc.)

[00229] According to one embodiment, each segment of the software may further include an indicator indicating whether the respective segment is a first, continue, or a last segment. According to one embodiment, the software includes one or more executable images that may be executed by repeater 2702 upon a completion of the download process. Thus, when repeater 2702 detects that a last segment of the software has arrived, repeater 2702 may start to execute the one or more executable images to construct an operating environment for operating repeater 2702. Alternatively, repeater 2702 may wait for a specific command from switch 2701 to start executing the executable images.

[00230] In one embodiment, each repeater includes a program that boots up the repeater when it is powered up and plugged into a port of a switch. The program may be an executable image stored in a bootable ROM or a flash memory of the repeater. When the repeater is powered up, the program is executed and starts to communicate with the switch. When the software or firmware is downloaded from the switch, the program loads the software or firmware into a memory, such as RAM of the repeater, and to optionally execute the software

or firmware to construct an operating environment (e.g., an embedded operating system) of the repeater for subsequent uses.

Alternative Exemplary Discovery Processes of Repeaters

[00231] Figure 27B is a flow diagram of another exemplary embodiment of a discovery process of a repeater. Exemplary process 2750 may be performed by one or more processing logic that may include hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both.

Referring to Figure 27B, during normal operations, switch 2751 periodically sends a heartbeat message to a network that includes one or more repeaters communicatively coupled to switch 2751 over the network (operation 2753). The heartbeat message sent by switch 2751 includes a configured VLAN ID associated with switch 2751 (e.g., VLAN tag field 2403 of Figure 24C having one of the aforementioned VLAN IDs), similar to exemplary switch heartbeat message 3450 of Figure 34B. When repeater 2752 is activated (e.g., powered up and plugged into one of the ports of switch 2751), it listens to all messages being transmitted over the network and specifically looks for a message of switch 2751 (operation 2755). Once repeater 2752 finds a message associated with switch 2751, repeater 2752 stores the VLAN ID of switch 2751 in a local memory (operation 2755) and starts to send heartbeat messages using the same VLAN ID. In response to the heartbeat messages received from repeater 2752, switch 2751 starts to download VLAN configuration information, including a VLAN ID for each type of transactions (e.g., exemplary VLAN configuration message 3700 of Figure 37), to repeater 2752 (operation 2758). The VLAN configuration information may be downloaded via one or more set data messages, similar to exemplary message 3600 of Figure 36A, as discussed above. Thereafter, repeater 2752 and switch 2751 may communicate with each other using an appropriate configured VLAN ID.

An Exemplary Software Download Process

[00232] Figure 28 is a flow diagram of an exemplary embodiment of a process for downloading software or firmware from a switch to a repeater using a tunneling protocol. Exemplary process 2800 may be performed by one or more processing logic that may include hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. Referring to Figure 28, in one embodiment, after a repeater has been configured using VLAN configuration information (e.g., VLAN configuration information message 3700 of Figure 37) downloaded from a switch, at block 2801, the repeater indicates to the switch that it is ready to receive software or firmware for setting up an operating environment from the switch. In one embodiment, the repeater may indicate via its heartbeat message, such as, for example, message 3400 of Figure 34A, by specifying the operating state, such as a “not loaded” state shown in Figure 35. Once the switch receives such indication, at block 2802, the switch starts to download the software or firmware to the repeater. In one embodiment, the software or firmware may be in an executable format and divided into multiple segments. Each segment may be downloaded to the repeater in sequence. The multiple segments of software or firmware may be downloaded using multiple set data messages, such as, for example, message 3600 of Figure 36A. In a particular embodiment, each segment may further include a memory location (e.g., the RAM within the repeater) where the respective segment should be loaded within the repeater.

[00233] Since the multiple segments are downloaded to the repeater in sequence, at block 2803, if the repeater determines that there is an error or the repeater misses at least one segment, the repeater may discard the software or firmware that has already been downloaded and reset itself, including a discovery process as described above (block 2806). Once all the software or firmware has been downloaded, the software or firmware is executed by the repeater to set up an operating environment of the repeater, similar to an operating system.

Once the operating environment has been set up, at block 2804, the repeater may enter into an operating state, such as, for example, an uninitialized state (see Figure 35), waiting for further instruction from the switch. In one embodiment, the repeater may report such operating state via a heartbeat message, such as, for example, heartbeat message 3400 of Figure 34A. Once the switch receives such message that indicates the repeater is in an idle state, at block 2805, the switch may configure the repeater to operate in certain states, such as, for example, a normal or a standby state. In one embodiment, the switch may configure the repeater via a set data message, such as, for example, exemplary set data message 3600 of Figure 36A.

An Exemplary Initial Token Assignment Process

[00234] Figure 29 is a flow diagram of an exemplary embodiment of a process for initial token assignment to a repeater. Exemplary process 2900 may be performed by one or more processing logic that may include hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. Exemplary process 2900 may be used when a mobile station enters a network having a switch and one or more repeaters, and the switch initially assigns a token to a repeater as a primary repeater communicating with the mobile station.

[00235] Referring to Figure 29, when mobile station 2901 enters the network, it broadcasts a message in the network, including switch 2903 and one or more repeaters 2902 (operation 2904), announcing that mobile station 2901 is entering the network. In one embodiment, the broadcasted message from mobile station 2901 includes at least one of a probe request and an authentication request conforming to the IEEE 802.11 specification. When the one or more repeaters 2902 receive the broadcasted message from mobile station 2901, repeaters 2902 forward the message to switch 2903 (operation 2905). In one embodiment, the broadcasted message from mobile station 2901 is encapsulated within an Ethernet packet via a tunneling protocol message, such as exemplary inbound 802.11 management message 3900 shown in

Figure 39A. Likewise, other 802.11 control and data messages may be forwarded from repeaters 2902 to switch 2903 using exemplary inbound 802.11 control message 3920 and inbound 802.11 data message 3940 shown in Figures 39B and 39C respectively. In the outbound direction, similar messages, such as 802.11 outbound management, control, and data messages, as shown in Figure 39D-39F, may be used. Particularly, the 802.11 outbound messages may be used by a switch to control the operating characteristics of a repeater regarding how to communicate with a specific client of the repeater, such as, for example, transmit power level, transmit rate, and transmit mode or channel, etc.

[00236] Referring back to Figure 29, once switch 2903 receives at least a portion of these messages from each of repeaters 2902, switch 2903 determines an appropriate repeater based on one or more policies, such as, for example, RSSI information of each repeater and assigns a token corresponding to mobile station 2901 to the selected repeater (operation 2906). Note that the RSSI information may be embedded in the respective 802.11 inbound message (see, exemplary 802.11 inbound messages shown in Figures 39A-39C. Alternatively, the RSSI information may be transmitted from each repeater to the switch via RSSI information messages (e.g., message 4200 of Figure 42). In one embodiment, switch 2903 assigns the token to the selected repeater via an assign token tunneling message. An exemplary message for assigning a token to a repeater is shown in Figure 40A. Likewise, a token may also be deleted or queried from one or more repeaters using exemplary token delete message 4020 and token query message 4040 shown in Figures 40B and 40C respectively.

[00237] Meanwhile, mobile station 2901 keeps trying to enter the network by continuously transmit 802.11 probe request or authentication request to the members of the network (operation 2907). Once the selected repeater receives the assigned token from switch 2903, the selected repeater starts to respond the messages (e.g., 802.11 probe request and authentication request messages) from mobile station 2901 (operation 2908) and thereafter, the selected repeater operates as a primary repeater (having the token assigned) to send

acknowledgement and to communicate with mobile station 2901 and to relay the 802.11 messages (e.g., management, control, data messages) between switch 2903 and mobile station 2901 (operation 2909).

An Exemplary Token Movement Process

[00238] After a token has been assigned to a first repeater, the token may be reassigned to a second repeater due to a variety reasons, such as, for example, poor signal quality based on RSSI or movement of the respective mobile station, etc. Figure 30A is a flow diagram of an exemplary embodiment of a process for token movement. Exemplary process 3000 may be performed by one or more processing logic that may include hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. Referring to Figure 30A, at block 3001, a mobile station detects that it is about to loose or has already lost a connection with the primary repeater it has been communicating based on, for example, errors, and starts to send several retry, probe request, and RTS (request to send) messages trying to restore the connection. At block 3002, these messages (e.g., retry, probe request, and RTS messages) are received by one or more repeaters and are forwarded to the switch via, for example, inbound 802.11 tunneling messages shown in Figure 39A-39C. The forwarded messages also include the RSSI information of each repeater performing the forwarding.

[00239] At block 3003, the switch receives these messages including the RSSI of each repeater and determines that there may be a “drop out” situation (e.g., either the mobile station exits the network or the repeater does not function correctly). The determination may be made based on one or more policies, such as, for example, the RSSI of the current primary repeater corresponding to the mobile station. If the switch determines that another repeater may be more appropriate to be the primary repeater, at block 3004, the switch performs a token handoff process that removes the token from the original primary repeater and assigns

the token to another repeater to enable that repeater as a new primary repeater. In one embodiment, the switch may transmit a token delete message to the original primary repeater to remove the token from that repeater using, for example, exemplary token deletion message 4020 shown in Figure 40B and assign the token to another repeater using, for example, exemplary token assignment message 4000 shown in Figure 40A. Other operations may be included.

Alternative Exemplary Token Movement Process

[00240] Figure 30B is a flow diagram of an alternative exemplary embodiment of a process for token movement. Exemplary process 3050 may be performed by one or more processing logic that may include hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. Referring to Figure 30B, at block 3051, each of the repeaters, including a primary repeater having a token for a mobile station, communicatively coupled to a switch periodically transmits its RSSI information (e.g., through a heartbeat message, such as message 3400 of Figure 34A or a data packet, such as 802.11 inbound messages shown in Figures 39A-39C) to the switch. Alternatively, the repeater transmits its RSSI to the switch when it is needed, such as, for example, when the RSSI drops below a certain threshold set by the switch. At block 3052, the switch keeps track the RSSI for each repeater and uses in part the RSSI information to determine whether a connection between the primary repeater and the mobile station is maintained at a reasonable quality level. At block 3053, if the switch determines that the connection may be about to drop (e.g., the RSSI of the primary drops below a threshold), at block 3054, the switch performs a token handoff process that removes the token from the original primary repeater and assigns the token to another repeater to enable that repeater becoming a new primary repeater for the mobile station, as described above.

Exemplary Link-Down Detection Process Performed By a Repeater

[00241] A link between a repeater and a switch may be broken due to a variety reasons, such as, for example, malfunction of the repeater, etc. As a result, one or more mobile stations communicating with the repeater may also loose the ability to communicate with the switch. Thus, it is important to be able to detect when such a condition occurs to avoid loosing the connection with the mobile stations. Figure 31A is a flow diagram of an exemplary embodiment of a process for detecting a link down. Exemplary process 3100 may be performed by one or more processing logic that may include hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. Referring to Figure 31A, at block 3101, a repeater determines that the link between the repeater and a switch may be down, based in part on the heartbeat, beacon, and/or data messages received from the switch. At block 3102, if the repeater determines that the link is down or about to go down, the repeater may perform a reset process by itself, including, but not limited to, a discovery, configuration, and downloading software or firmware from the switch again, using at least one of the aforementioned techniques, as described above.

Exemplary Link-Down Detection Process Performed By a Switch

[00242] In the above embodiment, the repeater that determines the link between the repeater and the switch may be down is still functioning. In such a case, all that is needed is for a reset process to bring the repeater back up. In some cases, the repeater may be dead or lost of power. In this situation, the switch needs to be able to detect that the connection with the down repeater is broken and to move the mobile stations communicating with the down repeater to another repeater. Figure 31B is a flow diagram of an alternative exemplary embodiment of a process for detecting link between a repeater and a switch. Exemplary process 3150 may be performed by one or more processing logic that may include hardware

(circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. Referring to Figure 31B, at block 3151, a switch determines that a link between the switch and a repeater may be down. The determination may be based on the heartbeat message (e.g., RSSI received from message 3400 of Figure 34A or no heartbeat message received at all), or other messages or indications (e.g., lack of response, etc.) received from the repeater. Alternatively, the determination may be based on the RSSI embedded within an inbound 802.11 message (e.g., messages shown in Figures 39A-39C).

[00243] At block 3152, the switch waits a predetermined period of time to allow the token of the suspected repeater to be “naturally” moved to another repeater (e.g., via a natural token movement algorithm), through the regular 802.11 retry, probe request, and RTS messages, which may be forwarded to the switch by other repeaters using one of the inbound 802.11 messages (e.g., messages shown in Figures 39A-39C). After the predetermined period of time expires, at block 3153, the switch determines whether there is still at least one mobile station “hung on” to the suspected repeater. In one embodiment, the determination may be performed based on an active station list stored in the switch. If so, at block 3154, the switch associates the one mobile station with another repeater. In one embodiment, the switch may transmit an 802.11 message, similar to a PING message, to the at least one mobile station via other repeaters to allow other repeaters to “hook up” with the mobile station. As a result, the mobile station is associated with another repeater, which becomes the primary repeater for the mobile station. Note that other processes, such as token assignment process as described above, may be invoked here. However, in one embodiment, if the at least one mobile station does not respond to such a message, the switch may consider that the mobile station has exited the network.

Exemplary Beacon Message Process

[00244] Beacon messages or beacon frames are part of 802.11 specification used to synchronize a switch, one or more repeaters, and the mobile stations. Figure 32 is a flow diagram of an exemplary embodiment of a process for handling a beacon message.

Exemplary process 3200 may be performed by one or more processing logic that may include hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. Referring to Figure 32, at block 3201, a switch periodically sends a beacon message (e.g., 802.11 beacon) to each of the repeaters communicatively coupled to the switch. Similar to other tunneling messages, the beacon message may be transmitted using a tunneling protocol. In one embodiment, the beacon message may be transmitted using an exemplary beacon frame 4100 shown in Figure 41. The data payload of the message contains 802.11 beacon.

[00245] In order to avoid collision of the beacon messages sent to each repeater, in one embodiment, the beacon messages are not transmitted at the same time. There may be a time offset between each beacon message. However, if the switch determines that substantial transmissions of the beacon messages would not cause interference that prevents the intended repeaters from receiving the beacon messages, the switch may transmit the beacon message at substantially the same time. In addition, according to one embodiment, the beacon message may include a time stamp performed by the switch. The time stamp represents an absolute time with respect to the switch.

[00246] At block 3202, when a repeater receives the beacon message from the switch, the repeater extracts the beacon message (e.g., 802.11 beacon message) from the respective Ethernet packet. In addition, the repeater may further perform a time stamp using its own timer and adding the timing information on the top of the time stamp of the switch. Since the transmissions of packets among the multiple repeaters communicatively coupled to the switch coordinate to function as an access point, each repeater may transmit or receive packets according to a schedule associated with the repeater. Particularly, when multiple repeaters are

operating at substantially the same frequency, the timing of the transmissions of packets for each repeater may be different. As a result, by the time that the repeater transmits the packet to a mobile station, the timing may be off due to delay. The mobile station receiving the packet readjusts the timing of its own to synchronize with the network. For example, according to one embodiment, when the repeater receives the beacon message, it may record the receiving time. Before the repeater transmits the packet to the receiving mobile station, the repeater may record the transmitting time and calculate the time difference between the receiving time and the time to transmit the packet to the mobile station (e.g., the time delta). The repeater then updates the time stamp by adding the time difference to the time stamp performed by the switch and transmits the packet to the mobile station (operation 3203). When the mobile station receives the packet, the mobile station can adjust its timing based on the time stamp performed by the switch and the repeater to synchronize with the network.

Exemplary Data Fragmentation Process

[00247] As described above, communications between a repeater and a switch are carried out via a tunneling protocol within the Ethernet protocol. That is, an ATP packet is encapsulated within an Ethernet packet transmitted between a repeater (e.g., repeater MAC sublayer) and a switch (e.g., switch MAC sublayer). However, as shown in Figures 24A-24D, an Ethernet packet itself includes an Ethernet header and the source and destination Ethernet addresses. An ATP packet also includes its own ATP header. As a result, in one embodiment, the data payload is limited to a maximum size of the corresponding interface between an ATP packet and the Ethernet packet less the Ethernet and ATP headers. When the data payload between the switch and repeater is large, for example, software to be downloaded or large 802.11 data message, a single Ethernet packet would not be able to carry such data payload. According to one embodiment, the data payload may be divided into multiple fragments when transmitted between a repeater and a switch.

[00248] Figure 33A is a flow diagram of an exemplary embodiment of a process for transmitting data payload in multiple fragments. Exemplary process 3300 may be performed by one or more processing logic that may include hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. Referring to Figure 33A, at block 3301, processing logic receives a data packet designated to a recipient, such as, a switch, a repeater, or a mobile station. At block 3302, processing logic divides the data payload of the packet into multiple fragments and encapsulates each fragment in a separate packet (e.g., one fragment per ATP packet embedded within an Ethernet packet). At block 3303, processing logic sequentially transmits the multiple fragments to the recipient. In one embodiment, each fragment includes an indication whether the fragment is a first, a continued, or a last fragment. In one embodiment, the multiple fragments of data are transmitted, wired or wirelessly, using a sequence of inbound or outbound 802.11 messages, similar to exemplary message shown in Figures 39A-39F. In one embodiment, each packet that carries a fragment may further includes a sequence number for identifying the current sequence, such that the recipient (e.g., a switch, a repeater, or a mobile station) can determine, based on the sequence number, whether one or more fragments have been missed.

[00249] On the receiving side, the recipient (e.g., a switch, a repeater, or a mobile station) may collect all fragments before processing the data. Figure 33B is a flow diagram of an exemplary embodiment of a process for receiving data payload in multiple fragments. Exemplary process 3350 may be performed by one or more processing logic that may include hardware (circuitry, dedicated logic, etc.), software (such as is run on a general purpose computer system or a dedicated machine), or a combination of both. Referring to Figure 33B, at block 3251, processing logic receives a packet (e.g., an Ethernet packet), including a data fragment of a data payload. The packet may be transmitted using a process similar to exemplary process 3300 of Figure 33A. At block 3352, processing logic determines whether

a previous fragment has been prematurely terminated, due a variety reasons, such as, for example, errors or missing one of the previous sequence. In one embodiment, processing logic simply determines based on the sequence number of the packet. If processing logic determines that there may be an abnormal termination (block 3353), processing logic may discard all the data that has already been received and restart over again. At block 3354, processing logic determines whether there are more fragments that have not been received yet. If so, processing logic may wait for a next fragment (block 3355). In one embodiment, each fragment includes an indication whether the fragment is a first, a continued, or a last fragment, such that processing logic can determine whether there are more fragments are expected. Once processing logic receives all fragments of the data payload, at block 3356, processing logic reassembles the data payload from the one or more data fragments received and processes the data payload as a whole. Other operation may be included.

[00250] Thus, tunneling protocols for wireless communications have been described herein. Whereas many alterations and modifications of the present invention will no doubt become apparent to a person of ordinary skill in the art after having read the foregoing description, it is to be understood that any particular embodiment shown and described by way of illustration is in no way intended to be considered limiting. Therefore, references to details of various embodiments are not intended to limit the scope of the claims which in themselves recite only those features regarded as essential to the invention.